There is now a published version of this paper ("The Discovery of Long-Run Causal Order: A Preliminary Investigation," Econometrics 8(3), 31, 3 August 2020, pp. 1-25: https://www.mdpi.com/2225-1146/8/3/31.)

Long-run Causal Order:
A Preliminary Investigation

Kevin D. Hoover*
Department of Economics
Department of Philosophy
Duke University

11 May 2018

*Address: Department of Economics, Duke University, Box 90097, Durham, NC 27708-0097, U.S.A. Telephone: +1 919-660-1876. Email: kd.hoover@duke.edu This paper arises out of a joint project with Søren Johansen, Katarina Juselius, and Morten Nyboe Tabor. I am grateful for many discussions and for comments on earlier drafts. A very early version of the paper was presented at the Econometrics Conference, Programme for Economic Modelling, University of Oxford, 1-2 September 2014. I thank the participants for valuable comments.

Long-run Causal Order: A Preliminary Investigation

Abstract

Graphical causal modeling has successfully provided algorithms for the empirical determination the contemporaneous causal order of structural vector autoregressions, but little attention has been paid to the long-run causal order of the cointegrated vector autoregression (CVAR) model. The main ideas of graphical causal modeling are reviewed. A distinction is drawn between ordinary variables that are nonstationary only when they have nonstationary causes and fundamental trends – i.e., variables whose own dynamics generate nonstationary behavior. The nature and limitations of causal relationships among fundamental trends is explored. It is then argued that the nonstationary behavior of CVAR models should typically be attributed to latent fundamental trends, so that their long-run dynamics arise from the causal structure connecting these fundamental trends to the ordinary variables and connecting the ordinary variables to each other. The connection between the graphical causal structure of the CVAR and weak exogeneity in different subsets of the CVAR variables is explored and some preliminary suggestions are offered on how to infer the underlying causal structure of the data-generating process through an exhaustive examination of weak exogeneity in irreducibly cointegrated subsets of the CVAR variables.

Keywords: graphical causal modeling, causal search, cointegrated vector autoregression (CVAR), weak exogeneity, irreducible cointegrating relations

JEL Classification: C32, C51, C18

Long-run Causal Order: A Preliminary Investigation

In the long run, we are all dead.
JOHN MAYNARD KEYNES

In the long run, we are simply in another short run.
VARIOUSLY ATTRIBUTED

Contrary to Keynes' famous dictum in the long run we are all dead,
the long run is with us every day of our lives
WALT ROSTOW

1. The Problem of Causal Order in the CVAR

Katarina Juselius's and Søren Johansen's most famous contributions to econometrics, studied and applied in detail in her textbook (Juselius 2006) and in his monograph (Johansen 1995) and in a large number of journal articles, concern the cointegrated vector autoregression (CVAR). The CVAR is a vector autoregression (VAR) in which special attention is given to the nonstationary components, or what we might consider to be the long-run properties, of the time series.

There are two significant traditions in time-series econometrics. The Cowles Commission in the 1940s and '50s pioneered *structural* econometrics that conceived of the econometric problem as one of articulating and measuring economic mechanisms (Koopmans 1950; Hood and Koopmans 1953; see Morgan 1990 for a history). The articulation of mechanisms was generally referred to as the *identification problem*. The major resource for securing identification was *a priori* economic theory. Early on, structural and causal articulation were regarded as synonymous, although subsequently causal language fell from favor (Hoover 2004). In his contribution to the 1953 Cowles volume, Herbert Simon (1953) drew on the language of experiments (actual or metaphorical) to suggest that an identified system of dynamic equations provided a map of the space of interventions in the economy.[1] Simon demonstrated an isomorphism between a structurally identified model and a causally well-ordered model.

A second econometric tradition, grounded more in time-series statistics, focused on *process* rather than structure (e.g., see Wold 1960 or Granger 1969). The VAR was introduced into macroeconometrics as part of a critical response to the Cowles Commission approach. Christopher Sims (1980, p. 1), building on earlier criticisms of Liu (1960) and others, attacked structural econometric models for making use of "incredible" identifying restrictions and offered the vector autoregression – a system of reduced-form equations in which all variables are endogenous – as a workable alternative to identified structural models. It rapidly became clear that reduced-form VARs were inadequate to counterfactual policy analysis – perhaps the most important use of macroeconometric models (Cooley and LeRoy 1985, Sims 1982, 1986). The *structural* VAR (*SVAR*), which imposes a causal order on the contemporaneous relationships among the endogenous variables was seen to provide the minimum restrictions needed to

[1] See Hoover (2001, ch. 3). In appealing to an experimental metaphor, Simon followed in the footsteps of Haavelmo (1944), a foundational figure for Cowles Commission Econometrics (see Hoover and Juselius 2015; Hoover 2014).

identify independent shocks, which were taken to be the drivers of a dynamic system, and policy analysis was largely reduced to working out the impulse responses to those shocks (see Duarte and Hoover 2012 and Hoover and Jordá 2001).

While the problem that had motivated Sims in the first place, the incredibility of the identifying restrictions, had been minimized, it was not eliminated, and the question, how we are to know the correct contemporaneous causal order, remains an open one. In truth, economic theory rarely provides a clear or decisive answer. In practice in most, though not all cases, SVARs were identified by assuming certain triangular causal orderings of the contemporaneous variables. Since all such causal orders are just-identified, they have the same likelihood function, and, thus, there is no empirical basis for choosing among them, so long as "empirical" is restricted to likelihood information. When the underlying data-generating processes (DGPs) are over-identified, information about conditional independencies among the variables provides information that can be used, in some cases, to distinguish among possible causal orders. This approach has been developed with great sophistication – mainly for non-time-series data – in the so-called graphical causality or Bayes-net literature (Spirtes, Glymour, and Scheines 2000; Pearl 2009). Swanson and Granger (1997) first applied a simple graphical causal search algorithm to the problem of determining the contemporaneous causal structure of an SVAR. Subsequently, more sophisticated algorithms have been applied and shown to be effective in a wide range of circumstances (Demiralp and Hoover 2003; Demiralp, Hoover, and Perez 2008).

Meanwhile, time-series econometrics discovered the importance of nonstationary processes and the concept of cointegration (Engle and Granger 1987). In light of these developments, the SVAR was reformulated into the CVAR. The issues that concern us, can be explicated in a standard CVAR with one lag and no deterministic components and variables integrated of degree one (notated I(1)):

(1) $$\Delta\mathbf{x}_t = \Gamma\Delta\mathbf{x}_{t-1} + \Pi\mathbf{x}_{t-1} + \boldsymbol{\varepsilon}_t,$$

where $\mathbf{x} = [x_1, x_2, \ldots x_p]'$, , $\boldsymbol{\Gamma}$ and $\boldsymbol{\Pi}$ are $p \times p$ matrices of parameters, $\boldsymbol{\varepsilon}_t \sim IN_p(\mathbf{0}, \boldsymbol{\Omega})$ is p-element vector of normal residuals, and t subscripts indicate time . The residuals contain both unobserved causes, which we shall call *shocks*, and various sorts of error. In general, the individual elements of $\boldsymbol{\varepsilon}$ are not independent, so that $\boldsymbol{\Omega}$ is not a diagonal matrix.

What Johansen and Juselius call a *structural* CVAR is derived from equation (1):

(2) $$\mathbf{A}_0\Delta\mathbf{x}_t = \mathbf{A}_1\Delta\mathbf{x}_{t-1} + \mathbf{a}\boldsymbol{\beta}'\mathbf{x}_{t-1} + \mathbf{u}_t$$

where $\mathbf{u}_t = \mathbf{A}_0\boldsymbol{\varepsilon}_t \sim IN_p(\mathbf{0}, \boldsymbol{\Sigma})$; $\mathbf{A}_0$ is a $p \times p$ with ones on the main diagonal such that $\boldsymbol{\Sigma} = \mathbf{A_0}\boldsymbol{\Omega}\mathbf{A_0'}$ and $\boldsymbol{\Sigma}$ is diagonal; and $\mathbf{A}_1 = \mathbf{A}_0\boldsymbol{\Gamma}$ (Juselius 2006, equation (12.2), p. 208; cf. equation (15.5), p. 276; Johansen 1995, pp. 78-79). If the variables in $\mathbf{x}$ are *cointegrated* (i.e., if a linear combination of nonstationary variables is itself stationary), then $\boldsymbol{\Pi}$ has reduced rank (r) and may be written as $\boldsymbol{\Pi} = \boldsymbol{\alpha\beta}'$, where $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ are $p \times r$ matrices, and $\mathbf{a} = \mathbf{A}_0\boldsymbol{\alpha}$.

Equation (2) is said to be *structural* because

1. the matrix $\mathbf{A}_0$, which is not in general unique, provided that it has $p \times (p-1)/2$ zero restrictions, identifies the causal ordering of the contemporaneous variables in the sense of Simon (1953; see also Hoover 2001, ch. 3);[2]
2. the long-run parameters of $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ can be recovered, provided that at least $r \times (r-1)$ restrictions are imposed on $\boldsymbol{\beta}$ (note, however, that in general $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ are not unique).
3. the elements of $\mathbf{u}$ (in contrast to those of $\boldsymbol{\varepsilon}$) are independent of each other and may, therefore, be taken to be the very shocks that are entangled in $\boldsymbol{\varepsilon}$, which is reflected in the fact that $\boldsymbol{\Sigma}$ is diagonal.[3]

Equation (2) is structural in, at best, a limited sense. One aspect of this claim can be grasped by comparing the practical importance of contemporaneous identification to the practical irrelevance of long-run identification. As noted $\mathbf{A}_0$ is not unique. Each admissible choice of $\mathbf{A}_0$ defines a distinct causal order and, in general, though all just-identified orders imply the same likelihood function, each will define a different set of shocks and different impulse-response functions. The matrices $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ are also in general not unique, but their product must equal $\boldsymbol{\Pi}$, and it is only $\boldsymbol{\Pi}$ that matters to the impulse-response functions. Thus, the choice among the admissible just-identified $\boldsymbol{\alpha}$s and $\boldsymbol{\beta}$s has no empirical consequences – at best, the choice is simply a matter of economic aesthetics.

There is a second reason to question the structural status of (2). If the rank of $\boldsymbol{\Pi}$ is $r < p$, then the variables in $\mathbf{x}$ are driven by $q = p - r$ common stochastic trends. It is possible that some of the common trends are embedded in the observable variables, but, generally, the common trends will be latent variables, which may, perhaps, be backed out of the observed variables (Juselius 2006, ch. 14). In that case, the essential elements of the structure, which are the causal drivers of its long-run behavior are not explicitly represented in (2), and (2) can, at best be a partial structural model

These two issues – the causal impotence of different choices of long-run identification and the latency of the principal causal drivers – are related. Together they form the major hurdle to applying graph-theoretic (so-called "Bayes-net") search algorithms to discover long-run causal structure.[4] The goal of this paper is to provide a coherent account of the causal order of a CVAR and to make some preliminary suggestions about how the methods of graphical causal search in conjunction with cointegration analysis might can aid in the *empirical* discovery of its long-run, as well as the short-run, causal structure.

2. Graph-Theoretic Causal Order

It will be helpful to review selectively some aspects of graphical causal analysis.[5]

[2] Since we are concerned principally about causal structure, we will focus on *generic* identification, setting to one side *empirical* and *economic* identification, in Johansen and Juselius's (1994) sense of these terms (see also Juselius's 2006, p. 208).
[3] The elements of **u** include both shock and error. Our focus here is on the shocks; in practice, we cannot ignore the error.
[4] We prefer the adjective "graph-theoretic" to "Bayes-net," as the search methods do not require a Bayesian approach to statistics.
[5] For compact treatments, see Cooper (1999) and Demiralp and Hoover (2003).

2.1 Graphs and Causal Structure[6]

In Simon's (1953) account, a structural model is a system of equations representing mechanisms in the world.[7] Although the account can be generalized considerably (see Hoover 2001, chapter 3), it will do for our purposes to restrict our attention to linear equations and to treat each equation as the representation for the causal mechanism determining its left-hand-side variable (the *effect*) in terms of right-hand-side variables (the direct *causes*). The coefficients on the right-hand-side variables are *parameters*, defined here to be variation-free variables that represent the loci of interventions in the system. That is, each parameter P_j, $j = 1, 2, \ldots n$, is defined as a set of possible values $P_j = \{p_{jk}\}$, where each p_{jk} may be an element in a discrete or continuous set. So, for example, a parameter might correspond to a switch or dial and its value to the particular setting of the switch or dial, and changing its value to an intervention in the system. The condition that parameters be *variation-free* means that any one parameter taking a particular value does not restrict the others from taking any value within their range (Hendry 1995, p. 163). We can define a *parameterization* (p_i) as the selection of one value for each parameter. Variation-freeness means, then, that the set of possible parameterizations of model is an element of the Cartesian product of the parameters: $p_i \in P_1 \times P_2 \times \ldots \times P_n$. A policy action that sets a parameter in, for example, a central-bank reaction function, is an example of an intervention. But parameters are not restricted to interventions by humans. Simon likens interventions to experiments and counts Nature among the experimenters. Indirect causal relations can be read off the system of equations from its recursive structure. Thus, if a variable (A) can be determined from a subset of equations of the system and another variable (B) can be determined only from a larger set of equations of which the first subset is an indispensable part, then A (directly or indirectly) causes B.

A system of equations that determines the values of a set of variables can be rewritten in various ways by taking linear combinations of its equations. If the true causal order corresponds to a particular way of writing the equations in which the coefficients on the right-hand-side variables are parameters as we have defined them, then any other linear combination does not represent the true causal order, as its coefficients are combinations (involving products, sums, and differences) of the parameters. These coefficients are not, therefore, variation-free, but involve cross-coefficient restrictions, and so are not themselves parameters. The problem of determining causal order, essentially the same as the classic identification problem, is to find a representation of the variables such that the interventions actually possible in the real world are represented as the parameters of the model. This is equivalent to finding a representation in which the coefficients on the right-hand-side variables are true parameters. Experiments (i.e., either actual human interventions in a real-world system or so-called "natural experiments") can provide the necessary evidence. Graph-theoretic search algorithms provide a method that can frequently derive the necessary evidence from passive observations.

To fix ideas, let us consider the causal relationships among a cross-section of variables with no time dimension. For example, in equation (1), consider a situation in which $\mathbf{\Gamma}$ and $\mathbf{\Pi}$ are identically zero. In that case, the causal structure among the variables in $\mathbf{x}$ can be represented by the matrix $\mathbf{A}_0$ in (2), where a non-zero element corresponds to the variable indexed by the row to the variable indexed by the column. In order to keep to our restriction to recursive systems, we

---

[6] This section borrows elements from Phiromswad and Hoover (2013), sections 2.1 and 2.2.

[7] Hoover (2001, chs. 2 and 3) provides a detailed account of Simon's approach and of it generalization to nonlinear systems, including ones with cross-equation restrictions among the parameters.

must consider only systems that are identified – that is, that impose at least $p \times (p - 1)/2$ zero restrictions (in addition to the restrictions imposed by the requirement that $\mathbf{\Sigma}$ be diagonal) – and that are lower triangular or can be made lower triangular simply by reordering the variables in $\mathbf{x}$.

The vector $\mathbf{x}$ is $p \times 1$, $\mathbf{A}_0$ is a $p \times p$ matrix of structural parameters a_{ij} with ones on the main diagonal, and $\boldsymbol{\varepsilon}$ a conformable vector of independent random shocks distributed $N(\mathbf{0}, \mathbf{\Sigma})$, where $\mathbf{\Sigma}$ is diagonal. Let the true DGP be a system of structural equations:

$$\mathbf{A}_0\mathbf{x} = \mathbf{u}. \tag{3}$$

Premultiplying (1) by $\mathbf{A}_0^{-1}$ yields the reduced form:

$$\mathbf{x} = \mathbf{A}_0^{-1}\mathbf{A}_0\mathbf{x} = \mathbf{A}_0^{-1}\mathbf{u} = \boldsymbol{\varepsilon}, \tag{4}$$

where $\boldsymbol{\varepsilon} \sim N(\mathbf{0}, \mathbf{\Omega})$ and $\mathbf{\Omega}$ is not in general diagonal.

The identification problem is the problem of working backwards from the observed probability model embodied in the reduced form (4) to the structural model (3). It is a problem because there are many matrices $\mathbf{P}^{-1}$, such that $E(\mathbf{P}\boldsymbol{\varepsilon}\,(\mathbf{P}\boldsymbol{\varepsilon})') = \mathbf{\Omega}$. Premultiplying (2) by $\mathbf{P}$ yields

$$\mathbf{P}\mathbf{x} = \mathbf{P}\mathbf{A}_0^{-1}\mathbf{u} = \mathbf{P}\boldsymbol{\varepsilon} = \mathbf{w}, \tag{5}$$

where $\mathbf{w} \sim N(\mathbf{0}, \boldsymbol{\varsigma})$ with $\boldsymbol{\varsigma}$ diagonal.[8] If (3) is the true DGP, the elements of $\mathbf{P}$ ($\neq \mathbf{A}$) are not structural parameters and the elements of $\mathbf{w}$ are not structural shocks, but pseudo-shocks that are complicated functions of the true shocks and the true structural parameters (a_{ij}s). Yet, because both (3) and (5) share the same reduced form (4), they cannot be distinguished from each other merely on their fit to observable data.

After normalization, the matrix $\mathbf{P}$ introduces $p(p - 1)$ additional parameters. Thus, any identification scheme must impose $p(p - 1)$ restrictions on the structural form. The fact that the covariance matrices of the shocks ($\mathbf{\Sigma}$ and $\boldsymbol{\varsigma}$) are diagonal imposes $p(p - 1)/2$ restrictions, so that identification requires that the coefficient matrices ($\mathbf{A}_0$ or $\mathbf{P}$) have a further $p(p - 1)/2$ restrictions. If there are exactly $p(p - 1)/2$ in each case, then the models are just identified and their likelihood functions are identical; so there is no choosing among them on the basis of the likelihood function. However, if there are more than $p(p - 1)/2$ restrictions, then the likelihood functions may be different and the overidentifying restrictions are testable. Causal search algorithms exploit those differences.

Equations (3) and (5) can be represented in distinct *directed graphs*. The elements of $\mathbf{x}$ form the *nodes* or *vertices* of the graph. An *edge* connects a pair of vertices. Edges come in several forms; for the moment, we consider just three: *null* or *no-edge*, indicating an absence of causal connection; an *unidirected* edge (—), indicating an undetermined causal direction; a *directed* edge, indicating an asymmetric causal influence ($\rightarrow$ or $\leftarrow$). If $a_{ij} \neq 0$, then a directed edge runs from variable j to variable i. A graph in which there are no undirected edges is called a directed graph.

Consider, for example, the system of equations (3) in which

[8] If $\mathbf{P}$ is lower triangular, then the transformation corresponds to a Choleski ordering familiar from structural VAR analysis.

$$\mathbf{x} = \begin{bmatrix} A \\ B \\ C \\ D \\ E \\ F \end{bmatrix} \text{ and } \mathbf{A}_0 = \begin{bmatrix} 1 & \alpha_{AB} & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ \alpha_{CA} & \alpha_{CB} & 1 & \alpha_{CD} & \alpha_{CE} & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & \alpha_{EF} \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}.$$

This system can be represented as in Figure 1. Each random shock in (1) could be treated as a latent variable and represented explicitly as connected by an edge directed into the dependent variable for its equation. To reduce clutter we take the random shocks as implied in the fact that the system is stochastic and would represent them graphically only if there were a special purpose in doing so.

It is sometimes useful to refer to the connections among the variables without reference to their directions. Every graph that connects the same variables is said to have the same *skeleton*. It is also useful to distinguish between a *parent*, the source of an arrow, and a *child*, the variable into which an arrow points. *D* and *E* are parents of *C*; *C* is a child of *D* and *E*. A variable that is the parent, grandparent, great grandparent, etc. of another variable is its *ancestor*. *F* is the ancestor of *C*. *B* is not an ancestor of *E*, because, even though *B* is connected to *C* and *C* to *E*, the arrows do not line up, so that causal influence is not transferred. By convention, every variable is regarded as its own ancestor. Variables that are connected by edges are said to lie on a *path*. When the edges align (e.g., on the path from *F* to *C*, then they lie on a *directed path*; whereas when the edges do not align (e.g., on the path *B* to *E*), they lie on an *undirected* path.

A graph is *acyclical* (which corresponds to what econometricians refer to as recursive systems) when descendants of any variable are not also ancestors of that same variable. Thus, Figure 1 is a directed acyclical graph (or *DAG*); while $A \rightarrow B \rightarrow C \rightarrow A$ is a directed *cyclical* graph. The graph $A \rightarrow B \rightarrow A$ is a very tight cyclical graph – namely, a simultaneous relationship between *A* and *B*, perhaps better represented as $A \rightleftarrows B$. Here, we will restrict ourselves to directed acyclical graphs – a typical *implicit* assumption in the VAR literature – even while recognizing that economists frequently invoke simultaneity.

2.2 Graphs and Conditional Independence

The key idea in graph-theoretic (or Bayes-net) accounts of causal structure is the mapping between the causal graph and the probability distribution described in the true DGP and its reduced form. The mapping is based on Reichenbach's (1956, p. 156) *principle of the common cause*: if any two variables, *A* and *B*, are probabilistically dependent, then either *A* causes *B* ($A \rightarrow B$) or *B* causes *A* or ($A \leftarrow B$) or they have a common cause ($A \leftarrow C \rightarrow B$). The principle of the common cause is generalized as the *causal Markov condition*:

Definition 1. *Let* G *be a causal graph relating a set of variables* ***V*** *with a probability distribution* P. *Let* ***W*** *be a subset of* ***V***. G *and* P *satisfy the* causal Markov condition *if, and* *only if, for every* ***W*** *in* ***V***, ***W****, conditional on its parents, is independent of every set of variables that does not contain its descendants*. (Spirtes *et al.* 2000, p. 29; see also Pearl 2009, p. 30).

Causal search is based in part on the systematic application of the causal Markov condition. Consider three simple cases: i) $A \leftarrow C \rightarrow B$; ii) $A \rightarrow C \rightarrow B$; and $A \leftarrow C \leftarrow B$. In each case, A and B are probabilistically dependent, but are independent conditional on C. The variable C in case i) is referred to as a *common cause* of A and B. In cases ii) and iii), C is an *intermediate variable.* In all three cases, C is said to *screen off A* from B.

Causal structure can in some cases induce conditional dependence between variables that are unconditionally independent or independent conditional on their parents. In Figure 1, D and E are unconditionally independent; but conditioning on C, their common effect, renders them probabilistically dependent. For example, let D = measure the dryness of the underbrush; E = the presence or absence of lightening; and C = the presence or absence of a forest fire. D and E may be completely independent. Yet, *if we know that the forest caught fire* and we know that the underbrush was dry, it raises the probability that the lightening was present. Vertex C in Figure 1 is called an *unshielded collider on the path DCE* (or *ECD*). It is a "collider" because the arrowheads come together at C, and is "unshielded" because there is no direct causal connection between C and E. Vertex C is a *shielded collider on the path ACB*; the edge $A \leftarrow B$ acts as a shield in that A and B are probabilistically dependent even without conditioning on the common effect.

Essentially, the causal Markov condition holds when a graph corresponds to the conditional independence relationships in the associated probability distribution. A graph is said to be *faithful* (by Spirtes *et al.* 2000, p. 31) or *stable* (by Pearl 2009, p. 31) if, and only if, there is a one-to-one mapping between the relationships of conditional independence implied by the causal Markov condition applied to G and those found in P (Spirtes *et al.* 2000, p. 48). There are well-known circumstances in which faithfulness can fail to obtain (see Sprites *et al*. 2000, p. 41; Pearl 2009, pp. 62-63; Hoover 2001, pp. 45-49, 151-153, 168-169). Essentially, faithfulness fails when the parameters are tuned in such a manner that variables that are in fact causally connected are nevertheless conditionally independent. Although failures of faithfulness can arise, they are not generic. Suppose, for example, that parameters were drawn out of uniform random distributions, then parameterizations that would violate faithfulness would typically have Lebesgue measure zero. There may be good economic reasons in some cases to expect the tuning that produces failures of faithfulness, but it is unlikely to arise by chance; and, in the spirit of Reichenbach's principle of the common cause, we should provide an economic account of how the fine tuning arises.[9]

The relationship between a causal graph and the probability distribution of the same variables is captured in Pearl's

[9] An example (see Hoover 2001, pp. 168-169): let $A = \varepsilon_A$, $B = \alpha A + \varepsilon_B$, and $C = \beta A + \gamma B + \varepsilon_C$, where the Roman letters indicate variables, the εs indicate mutually independent random error terms, and other Greek letters indicate parameters. Generically, in this system $A \rightarrow C \leftarrow B$. However, if the parameters happen to fulfill the condition $\alpha = -\beta/\gamma$, then the conditional correlations corr($AC|B$) = 0 and corr($BC|A$) = 0 – that is, there are independence relationships that are not encoded in the graph, which is a failure of faithfulness. While it is unlikely to arise by chance, this very special parameterization may occur nonetheless for good economic reasons: it is the condition that would be fulfilled by a policymaker trying to use A optimally to minimize the variance of C.

d-separation Theorem. *Let* G *be a causal graph relating a set of variables* ***V****, and let* ***X****,* ***Y****, and* ***Z*** *be distinct subsets of* ***V****.* ***X*** *and* ***Y*** *are* d-separated *given* ***Z*** *and, thus, probabilistically independent conditional on* ***Z*** *if, and only if, there exists no undirected path* U *between* ***X*** *and* ***Y*** *such that i) every collider on* U *has a descendent in* ***Z****, and ii) no other variable on* U *is in* ***Z****; otherwise,* ***X*** *and* ***Y*** *are* d-connected (Pearl 2009, pp. 16-18; Spirtes *et al.* 2001, p. 44).

Some examples of *d*-separation in Figure 1:

a) *D* and *E* are d-separated conditional on the null set, since there is only one path between them, which is blocked by the collider *C*; they are, therefore, unconditionally independent;

b) *D* and *E* are not d-separated conditional on *C*, since by clause i) the collider *C* is a member of the conditioning set (recall that a variable is considered to be its own descendant), which corresponds to what we already learned – namely, that conditioning on a unshielded collider renders unconditionally independent variables conditionally probabilistically dependent;

c) *C* and *F* are d-separated conditional on *E*, since there are no colliders on the only path between them, so that clause i) does not apply, and there are no other variables on the path, so that clause ii) does not apply, which corresponds once again to what we have already learned – namely, that an intermediate variable on a directed path screens off probabilistic dependence, provided that there are no alternate paths.

The d-separation theorem allows us to read from a causal graph which conditional independence relationships are encoded in the likelihood function.

Different graphs may imply the same set of conditional independence relationships, so that the corresponding probability distribution defines a class of observationally equivalent causal structures. This class may have only one element or it may have many. The class of admissible orderings is characterized by the

Observational Equivalence Theorem. *Any probability distribution that can be faithfully represented in an acyclical graph can equally well be represented by any other acyclical graph that has the same skeleton and the same unshielded colliders* (Pearl 2009, p. 19, Theorem 1.2.8; see also Sprites *et al.* 2000, ch. 4).

The theorem implies that there may be causal structures in which some causal edges are reversed and yet all of the unshielded colliders preserved. For example, in Figure 1, reversing the edge $F \rightarrow E$ leaves the likelihood unaffected, so that information about the probability distribution of the variables cannot by itself provide a basis for inferring the causal order of that edge in the true DGP. A just-identified model has no unshielded colliders. It follows immediately that all just-identified models of the same variables are observationally equivalent. Consequently, all Choleski orderings of the contemporaneous variables in an SVAR are observationally equivalent.

2.3 Causal Inference in Causally Sufficient Sets

The existence of observationally equivalent causal structures implies that the graphs of such structures may not be uniquely recoverable from information in embedded in the likelihood function. The best case occurs when the data are complete in the special sense of being causally sufficient:

Definition 2. *A set of variables is* causally sufficient *if, and only if, any variable that is excluded from the set directly causes at most one variable within the set* (see Spirtes *et al.* 2000, p. 22).

Consider a simple example. Suppose that a DGP includes just three variables, $A \leftarrow L \rightarrow B$. The complete set of variables is causally sufficient, and the fact that *A* and *B* are independent conditional on *L* would allow us to infer $A — L — B$, although it would not allow us to recover the DGP graph. If only *A* and *B* were observed (i.e., *L* were latent) and we wrongly assumed that they formed a causally sufficient pair, then we would, incorrectly, find that they were not independent and would treat them as directly connected, $A — B$. Notice that causal sufficiency implies that shock terms, which are unobserved, latent random variables must be independent, as they may cause at most one of the observed variables.

There is a variety of causal search algorithms for causally sufficient acyclical directed graphs.[10] It is unnecessary to describe them in detail, although a schematic account of two popular ones – the PC and SGS algorithms, which differ only in ways that are unimportant for our purposes – will help to fix ideas.[11] Each consists of three main steps: 1) starting with a set of variables assumed to be causally sufficient and densely connected with undirected edges, test all possible conditional independence relations and eliminate edges whenever any pair of variables is unconditionally or conditionally independent; 2) identify unshielded colliders (i.e., for every triple of variables in which two of the variables are independent conditional on some set of variables but not conditionally independent when a particular variable is added to the conditioning set) and orient the causal arrows towards this last conditioning variable; and 3) apply logical constraints to orient as many additional causal arrows as possible – in particular, orient arrows to avoid creating unshielded colliders unsupported by the conditional independence tests in step 3) and to avoid creating cycles.

Some search algorithms work on other principles or allow cyclical causal orders or latent variables. Search algorithms may also be modified to incorporate *a priori* knowledge, either by insisting on, or forbidding, particular causal connections irrespective of the statistical information. Time structure may, for example, be imposed as *a priori* knowledge.

2.4 Discovering Causal Structure with Latent Variables

Since a search algorithm cannot work as described in Section 2.3 if there are latent variables that violate causal sufficiency, other algorithms have been developed that extract additional information from the observed variable set about possible causal patterns between latent and non-latent variables (e.g., Spirtes *et al.*'s (2000, chs. 10 and 11), FCI algorithm). Some of these search algorithms return *partial ancestral graphs* as their output. In addition to unidirectional arrowheads, partial ancestral graphs can connect observable variables with a bidirectional edge ($\leftrightarrow$) or an edge with an open circle ($\circ$) as one or both of its endpoints. The bidirectional edge is interpreted not, it must be emphasized, as a simultaneous relationship, but rather as indicating that two variables have a latent common cause. Thus, $A \leftrightarrow B$ is interpreted to mean

[10] See Cooper (1999), Spirtes *et al.* (2000, chs. 5 and 6), and Pearl (2009), ch. 2. The *Tetrad* software package implements Spirtes *et al.*'s (2000) algorithms2, as well as additional algorithms available on Carnegie-Mellon University's Tetrad Project website: http://www.phil.cmu.edu/projects/tetrad/.
[11] The first is named for the first initials of its creators: Peter Spirtes, and Clark Glymour; the second for initials the surnames of its creators: Spirtes, Glymour, and (Richard) Scheines.

$A \leftarrow L \rightarrow B$, where L is an unobserved latent variable. The open circle is interpreted as *either* an arrowhead or as a blank endpoint. Thus, $\circ\!\!\rightarrow$ represents that the data are consistent with a DGP in which either $\rightarrow$ or $\leftrightarrow$ is the true causal edge.

3. The CVAR as a Nearly Decomposable System

3.1 Near Decomposability in Nonstationary Systems

It is well known in the CVAR literature, that there is a substantial statistical independence between $\Delta\mathbf{x}$, reflecting short-run or stationary properties of the variables and the lagged levels of $\mathbf{x}$, reflecting the long-run or nonstationary behavior (Juselius 2006, p. 208). The independence appears to be closely related to Herbert Simon's (1996, ch. 8) analysis of hierarchies of complex systems based on the notion of near decomposability.

A system is *decomposable* in Simon's account when some parts are tightly related to one another so that they form a unit, and the unit interacts with other units only as a whole. Decomposability is rare, yet near decomposability is more common. A system for Simon is *nearly decomposable* when

> (1) . . . the short-run behavior of the component subsystems is approximately independent of the short-run behavior of the other components; [and] (2) in the long run the behavior of any one of the components depends in only an aggregate way on the behavior of the other components. [Simon 1996, p. 198; see also Boumans 2005**,** ch. 4 and Hoover 2015, and, for formal treatments, Simon and Ando 1961 and Simon and Iwasaki 1988.]

Simon's idea is that even when systems are densely causally connected, some linkages are strong and/or fast-acting, while others are weak and/or slow-acting (see also Simon and Rescher 1966). The variables that are connected by strong/fast linkages act as units that can be analyzed independently and, subsequently, can be analyzed in relationship to similar units.

To take an economic example, on a quarter-to-quarter or year-to-year basis, the level of employment in an economy is dominated by a complex interaction among relatively quickly adjusting factors, such as wage rates, interest rates, private and government expenditure. Yet, on a century-to-century basis, the level of employment is dominated by the slowly changing level of overall population, which probably has little to do with the detailed factors that govern employment in the short run, though it may well have to do with the also relatively slowly changing level of technology. Similarly, the rates of return on various financial assets on a minute-to-minute or day-to-day basis are determined mainly by the rapid process of arbitrage in financial markets, while the aggregate level of all interest rates on a year-to-year basis is determined by the more slowly developing profitability of industry in general.

Simon provides no criterion for distinguishing fast and slow adjustment. Our two economic examples, in fact, suggest that it will often be a matter of degree or a matter of how the analysis is contextualized. In the case of simple dynamic processes, the speed of adjustment may be captured by the value of the autoregressive parameter. Simon does not suggest that the distinction between fast and slow is sharp. In the case of processes such as equation (1), a natural division point is the difference between stationary and nonstationary processes – i.e., between processes with and without unit roots – since there is a sharp distinction between the probability models appropriate in each case. The transition between the two families of probability distributions represents a kind of critical phenomenon familiar in certain physical applications (e.g., see Batterman 2001). There is a sense, then, in which the long-run causal structure emerges from the short-run structure at the point of transition between stationary and nonstationary processes. The short-run and long-run causal structures may, in fact, pull in opposite directions.

The CVAR also illustrates another feature of near decomposability – temporal aggregation. Short-run dynamics involve all the variables in the system, while, in the long-run, groups of variables are dominated by a fewer number of common trends. This explains cointegrating relationships are robust to widening the data set and recommends a *specific-to-general* approach: once the trends can be characterized, then any new variable is either redundant or carries information with respect to a new trend (Juselius 2006, ch. 22; Johansen and Juselius 2014). The minimum number of variables needed to characterize a trend corresponds to Davidson's (1998) *irreducible cointegrating vector*; but, since any set of variables adequate to characterize the trends is as good as any other, irreducible cointegrating relationships will not, in general, be unique.

The "near" in near decomposability reminds us that the short-run and long-run behaviors of the system are not completely disconnected. The underlying trends must be transmitted to variables that would not trend on their own through short-run adjustments. This is the meaning of Rostow's epigraph "the long-run is with us every day of our lives." These short-run adjustments ultimately account for the cointegration of the nontrend variables with the trend and with each other. The implication in the context of the CVAR is that a meaningful empirical identification of the long-run structure of the nontrend variables, must address the manner in which the long-term elements affect short-run behavior. There are conceivably causal relationships directly among trends, but all other causal relationships must be mediated through stationary behavior.

3.2 The Existence of Nonstationary Trends

Although this last point requires further amplification, we must first consider the status of the nonstationary trends themselves. The variables in a system such as equation (1) may all be I(1) and cointegrated when some of the eigenvalues of $\Pi = 0$ and, equivalently, when some of the rows of Π are linearly dependent, resulting in reduced rank. These conditions may be fulfilled by a precise tuning of the parameters of Π. In such a case, despite the fact that the CVAR literature refers to the cointegrated variables as being driven by common stochastic *trends*, there are, in fact, no actual trends, in the sense that there are no observed or latent variables that are the trends; rather, the trends are virtual only: they are special (and non-robust) properties of the causal systems and not ontologically independent mechanisms distinct from the other variables in the system. We can, of course, express equation (1) in a moving-average form in which each variable is determined by a weighted average of q ($= p - r$) shock processes, where p is the number of variables and r is the rank of Π. But these shock processes are, in fact, themselves simply weighted averages of the shocks to individual variables (see Johansen 1995, chs. 3-4; Juselius 2006, chs. 14-15). In this case, the common-trends representation does not correspond to anything in the world, but is simply a convenient repackaging of the information in (1).

Such virtual common stochastic trends are not robust in the sense that deviations from the precise parameterizations of Π that generated them will typically result in a collapse of the cointegration among the variables in **x** and the disappearance of the common-trends representation and, very possibly, result in some of the variables turning out to be I(0) rather than I(1). In this case, we can say that the common trends are not *generic*. If we imagine that the parameters of $\boldsymbol{\Pi}$ are drawn from continuous uniform distributions, then the common trends are not generic in the same way that the failures of faithfulness discussed in Section 2.2 (esp. fn. 10) are not generic (cf. Davidson 1998. pp. 90). As with failures of faithfulness, there may well be cases in which virtual trends can be expected to appear for good economic reasons because precise tuning of the $\boldsymbol{\Pi}$ parameters. But absent such specific cases, our concern shall be exclusively with generic trends – that is, trends that genuinely exist in the world. (These issues are discussed in detail from a philosophical point of view in Hoover 2015.)

When can we say that trends are genuine features of the world? A trend might be observable – that is, there is a variable in the model that corresponds to a mechanism in the world that generates I(1) behavior and transmits it to other variables, where this behavior does not depend on the parameterization of the other variables in the system. Some combinations of the other variables that are causally connected to these trend variables will be cointegrated robustly with respect to reparameterization of other relationships in the system.

A trend might be real (i.e., not virtual) and yet be unobservable or unobserved (i.e., latent). To alter the story a little, imagine that there are trends causally connected to the variables in **x** but they are unobserved, latent variables and, so, not included in **x**. In that case, in general, as in other cases of omitted variables, some of the coefficients on some of the variables in **x** will not be variation-free, but will have dependencies, including cointegrating relations, dictated by the omitted causal structure. The coefficients are not, then, true (structural) parameters.

How can we distinguish a virtual trend from a genuine unobserved, latent trend? Johansen and Juselius (2014) show that when trends are genuine we may obtain essentially the same estimates for them from narrower and wider sets of cointegrating variables. The precise values of estimates of the common stochastic trend and the trend shock will be different, but asymptotically, as the number of observations becomes large, the estimates from the narrower and wider data set will converge. While for any one set of data we cannot rule out that the parameters are perfectly tuned to produce cointegration, it would generally require a different tuning to produce the same result in a wider data set. Thus, the stability of the estimates of the trend and its shock across increasingly wider data sets is strong evidence that the trend is, in fact, a genuine, ontologically independent, though not directly observed, variable.

3.3 Causal Order in the Long Run

3.3.1 *Fundamental Trends*

What exactly is meant by long-run causal order? Consider a simple structural dynamic system:

(6) $$x_t = \mu + \alpha x_{t-1} + \beta(y_{t-1} - y_{t-2}) + \varepsilon_t^x$$

(7) $$y_t = \gamma + \delta y_{t-1} + \lambda x_{t-1} + \varepsilon_t^y .$$

This system can be represented in the causal graph in Figure 2.A.

The long run at time *t* can be defined to be the counterfactual situation towards which the system would be heading if it were subjected to no further shocks (i.e., when $\varepsilon_{t+n}^{x} = \varepsilon_{t+n}^{y} = 0$, for all $n \geq 1$) and the dynamics are allowed to fully work themselves out, so that the variables are no longer changing (i.e., $x_{t+n} = x_{t+n-1}$ and $y_{t+n} = y_{t+n-1}$ as $n \to \infty$). Call these long-run values of the variables x_t^{∞} and y_t^{∞}. They remain indexed by *t*, since where the variables will end up depends on the point at which we stop shocking the system.

Applying these ideas, each of the equations (6) and (7) can be transformed to its long-run version by setting all shocks and all differenced variables to zero, and all level variables to their common values, and simplifying equation by equation (in order not to disturb the structural nature of the equations). Consider first the case in which $|\alpha| < 1$ and $|\delta| < 1$, so that the variables are stationary:

(8) $$x_t^{\infty} = \frac{\mu}{1-\alpha}$$

(9) $$y_t^{\infty} = \frac{\gamma}{1-\delta} + \frac{\lambda}{1-\delta} x_t^{\infty}.$$

The causal graph of this long-run structure is given in Figure 2.B, which is much less complex than that of Figure 2.A. Also, note that, while *x* and *y* (at different lags) cause each other in the system (7) and (8) and Figure 2.B, *x* is a one-way cause of *y* in the long-run.

In the stationary case, the system settles down to fixed values, regardless of initial conditions. In contrast consider a nonstationary case in which $\mu = \gamma = 0$ and $\alpha = 1$. Under this parameterization, the long-run system is then

(10) $$x_t^{\infty} = x_t^{\infty}$$

(11) $$y_t^{\infty} = \frac{\lambda}{1-\delta} x_t^{\infty}.$$

Although equation (10) shows that x_t^{∞} is not caused in the long-run by y_t^{∞}, unlike in the stationary case, it has its own dynamic, and the value of x_t^{∞} changes as time advances, so that we can write (10) as

(10)′ $$x_t^{\infty} = x_t.$$

We take up this point in the next section.

The long-term causal relation in (11) is still appropriately represented in Figure 2.B. It is important to note that the equations (10) or (10)′and (11) (and, equally, equations (8) and (9)) do not represent the reduced-forms or *explicitly* the values of the variables in the long run any more than equations (6) and (7) represent explicitly the value of the variables along the dynamic paths. Rather they are the long-run causal structural equations (causes on the right, effects on the left), which, of course, can be solved to find the long-run values. Figure 2.A shows that if we intervened in the process governing *y* (e.g., by changing one of the parameters λ or δ), it would have various dynamic implications for *x*, as well as for *y*. But Figure 2.B and its associated equations (10)′ and (11) show that, in the long-run, x is ordered recursively ahead of y, so that the intervention would not transmit from y to x, once all the dynamics had been worked out.

We would not, therefore, in this structural form substitute (10)′ into (11), to give the long-run solution for y_t^{∞}, as such a substitution and, in general, taking linear combinations of distinct structural equations, as one does in solving for a reduced-form, destroys the correspondence between the systems of equations and the causal structure, represented, for example, in a causal graph. We now turn to characterizing the sources of long-run trends, such as x_t^{∞} and of their causal relationships to each other and to non-trend variables.

3.4 Fundamental Trends and Ordinary Variables

The order of integration of time series is not a structural characteristic but the phenomenal consequence of the structural characteristics of the data-generating process. Consider a simple I(1) process:

(12) $$x_t = x_{t-1} + \varepsilon_t ,$$

where time t = 0, 1, 2, . . . The process can be expressed in moving-average form:

(13) $$x_t = x_0 + \sum_{j=1}^{t} \varepsilon_j ,$$

where x_0 is an initial value. Similarly, a simple I(2) process

(14) $$\Delta x_t = \Delta x_{t-1} + \varepsilon_t ,$$

can be expressed as

(15) $$x_t = x_1 + \Delta x_1 (t-1) + \sum_{i=2}^{t} \sum_{j=2}^{i} \varepsilon_j ,$$

where Δx_1 and x_1 are initial values.

Now, consider a variable y, which is I(1) but has an independent I(1) cause z. Thus, z is described by analogues to (12) and (13):

(16) $$z_t = z_{t-1} + \varepsilon_t^z ,$$

and

(17) $$z_t = z_0 + \sum_{j=1}^{t} \varepsilon_j^z .$$

And y can be expressed as

(18) $$y_t = y_{t-1} + \alpha z_{t-1} + \varepsilon_t^y ,$$

where α is a parameter measuring the strength of the causal connection between z_t and y_t. Substituting (17) into (18), yields

(19) $$y_t = y_{t-1} + \alpha\left(z_0 + \sum_{j=1}^{t-1} \varepsilon_j^z \right) + \varepsilon_t^y .$$

The moving-average form of (19) is

(20) $$y_t = y_1 + \alpha z_0 (t-1) + \sum_{j=2}^{t} \varepsilon_j^y + \alpha \sum_{i=2}^{t} \sum_{j=2}^{i} \varepsilon_{j-1}^z ,$$

where y_0 and z_0 are initial values.

Equation (20) shows that the variable *y* is I(2). But consider the thought-experiment in which $\varepsilon_t^z = 0$ for every *t*, which is the same as the causal connection between *z* and *y* having been severed. In that case, *y* – now driven entirely by its internal dynamic – would be I(1) and not I(2). Looked at as a "dynamic processor," the structural equation (12) for *y* is fundamentally an I(1) generator. The variable *y* becomes I(2) only because the I(1) series *z* is run through this I(1) generator, which raises its natural order of integration by one degree.

A key feature is captured in a key distinction between ordinary and own orders of integration:

Definition 3. *The* ordinary order of integration *of a variable* x *is the number of times it must be differenced to render it stationary – that is, it is the property than we have hitherto indicated by saying x is I(*n*), which says that* Δ^nx *is stationary and* Δ^{n-1}x *is not*.

Indicate the ordinary order of integration for a variable *x* as I(*x*), so that "*x* is I(*n*)" can also be written as I(*x*) = *n*. When we use "order of integration" without qualification, we mean the ordinary order of integration.

Definition 4. *The* own order of integration *of a variable* x *(indicated by the operator Ξ(*x*)) is the order of integration that would result from the structural equation for x, considered independently from all other structural equations and setting all variables, except for* x*, its lagged values, and its own shock, to zero.*

Thus, for equations (12) and (16), Ξ(*x*) = Ξ(*z*) = 1 and I(*x*) = I(*z*) =1; and for equation (14), Ξ(*x*) = 2 and I(*x*) =2; whereas for (18), the own and ordinary orders of integration diverge, so that Ξ(*y*) = 1 and I(*y*) = 2.

It is easy to show that if a variable *y* has causes z_1, z_2, . . . z_n, the relationship of the ordinary to the own order of integration is given by

(21) I(*y*) = Ξ(*y*) + max[I(z_1), I(z_2), . . . I(z_n)].

We can now define explicitly the notions that we used implicitly earlier:

Definition 5. *An* ordinary variable *(*x*) is a variable for which Ξ(*x*) = 0.*

Definition 6. *A* nonstationary processor *(y) is a variable for which Ξ(*y*) ≥ 1.*

Definition 7. *A* fundamental trend *(*T*) is the nonstationary component of a nonstationary processor (*y*) for which Ξ(*y*) = I(*y*).*

We can demonstrate:

Proposition 1. *The nonstationary components of fundamental trends are probabilistically independent of each other*.

Proof. Consider two fundamental trends, x, for which $\Xi(x) = I(x) = m$ and y for which $\Xi(y) = I(y) = n$. Assume that they are probabilistically dependent. By Reichenbach's Principle of the Common Cause, variables are probabilistically dependent if, and only if, one is the cause of the other or both are the effects of a common cause. Consider the case in which one is the direct cause of the other: $x \rightarrow y$. Since y is a nonstationary processor, its order of integration should be $I(y) = \Xi(y) + I(x) = m + n$; but that contradicts the assumption that $\Xi(y) = I(y) = n$ and, therefore that y is a fundamental trend. It follows, then, that *not* $(x \rightarrow y)$, which in turn implies that probabilistic dependence cannot be induced by a direct causal connection when the effect is a fundamental trend. Consider the second case in which x and y have a common cause z: $x \leftarrow z \rightarrow y$. If $\Xi(z) > 0$, a similar argument shows that neither x nor y could be a fundamental trend. And if $\Xi(z) = 0$, then the probabilistic dependence induced by z between x and y is confined to the stationary components of those variables. Thus, either we must reject that z is a common causes of x and y and, therefore, that x and y are probabilistically dependent or that the probabilistic dependence involves the nonstationary components. We can conclude, therefore, that x and y cannot be fundamental trends and have probabilistically dependent nonstationary components.

When one variable causes another, the fundamental trends of the cause are passed on to the effect. An ordinary variable can be I(1), so long as the largest order of integration of one of its causes is $I(z_j) = 1$. Equation (15) shows that that there could be more than one I(1) cause, since the order of integration depends on the maximum order of integration. Thus, an ordinary variable y can be caused by any number of I(1) variables and yet remains I(1). On the other hand, if $\Xi(y) = 1$ and the maximum ordinary order of integration among its causes were $I(z_j) =1$, then $I(y) = 2$.

We can define the element that accounts for the cointegration among variables:

Definition 8. *A* local trend *(*T*) is the linear combination of fundamental trends that constitutes the nonstationary component of a variable.*

Since a linear combination can place a weight of zero on one of its constituents, a fundamental trend is trivially a linear combination of itself and any other fundamental trends with zero weights. Consequently, a fundamental trend is also a (degenerate) local trend.

If two variables are cointegrated, then they share a common local trend.

The notion of own order of integration and the distinction between ordinary variables and fundamental trends combined with some commonplace empirical observations allows us to say some relatively important things about long-term causal structure.

Consider the causal relations among variables for which the largest ordinary order of integration is I(1). These variables are typically analyzed as being driven by stochastic trends fewer in number than the total number of variables. In such a system, the following propositions are true:

Proposition 2. *The fundamental trends must necessarily be I(1).*

Proposition 3. *Any variable* x *in the system that is not identical with one of the trends has* $\Xi(\mathrm{x}) = 0$.

Proposition 4. *Fundamental trends can cause* only *ordinary variables and not other fundamental trends.*

Proof: Any fundamental trend T_j in the system has the property of being both $\mathrm{I}(T_j) = 1$ and $\Xi(T_j) = 1$; but if any such trend had a cause that were I(1), then it would have to be that $\mathrm{I}(T_j) = 2$, which is a contradiction.

Turning to systems that contain both I(1) and I(2) variables. In practice, econometricians rarely find more than one I(2) variables in a data set. Restricting ourselves to the case of a *single* I(2) variable:

Proposition 5. *The I(2) variable can cause only ordinary variables and not another fundamental trend.*

Proof. Analogous to the argument for proposition 4: if it did, it would generate an I(3) variable, which *ex hypothesi* is not in the system.

Proposition 6. *The own-order of integration of the I(2) variable* (x) *is either* $\Xi(\mathrm{x}) = 2$, *in which case, its nonstationary component is a fundamental trend, or* $\Xi(\mathrm{x}) = 1$ *with one or more I(1) trends causing it, in which case all the fundamental trends in the system* ($\mathrm{T_j}$) *are* $\Xi(\mathrm{T_j}) = 1$.

Although these six propositions very likely do not exhaust what can be learned about the causal structure of long-run relations from this fundamentally simple analysis, two points are striking and, perhaps, somewhat surprising: first, unlike the case of short-run relationships, substantial conclusions about causal order may be inferred from facts about integration without any need to appeal to an intransitive conditioning relationship; second, the analysis combined with some commonplace empirical observations known to most practitioners of CVAR analysis, suggest that the causal structure of the long-run is simpler – there are fewer and less dense causal connections among trends than among ordinary variables.[12]

4. Graphical Analysis of the CVAR

4.1. The Canonical CVAR of a Causally Sufficient, Acyclical Graph

Consider first the long-term structure of a causally sufficient CVAR with an acyclical causal structure in which the fundamental trends are represented explicitly. In the remainder of the paper, we consider only cases for a strong form acyclicality in which we do not permit any feedback from one variable to another, even with a time delay. Thus, we rule out cases such as $X_t \rightarrow Y_{t+1} \rightarrow X_{t+2}$. Consider only the case in which all variables are I(1) and in which the I(0) dynamics have been concentrated out and in which the contemporaneous causal order has been imposed:

(22) $$\Delta \boldsymbol{\xi}_t = \boldsymbol{\Psi}\boldsymbol{\xi}_{t-1} + \mathbf{H}_t ,$$

[12] Note, however, that we have continued to restrict consideration to recursive orderings and not allowed for simultaneity or cyclicality. We conjecture, but do not provide the analysis here, that if both $\Xi(T_h) = 1$ and $\Xi(T_k) = 1$ and the two trends are simultaneous ($T_h \leftrightarrows T_k$) both series will prove to be I(2).

where $\boldsymbol{\xi} = [\mathbf{X}', \mathbf{T}']'$; **T** is a $q \times 1$ vector of fundamental trends; **X** is a $p \times 1$vector of ordinary variables, which may be trending (i.e., I(1)), but are not fundamental trends; $\mathbf{H}'_t = [\varepsilon_{1,t}, ..., \varepsilon_{p,t}, \eta_{1,t}, ..., \eta_{q,t}]'$ is a $(p+q) \times 1$ vector of shocks to ordinary variables (ε_{it}, $t = 1, 2, \ldots, p$) and to fundamental trends (η_{jt}, $j = 1, 2, \ldots, q$), each of the elements of which is an identically independently distributed random variable and $\mathbf{H}_t \sim IN(\mathbf{0}, \Omega)$, where $\boldsymbol{\Omega}$ is diagonal.

The system can be partitioned as

(23) $$\Delta\boldsymbol{\xi}_t = \begin{bmatrix} \Delta\mathbf{X} \\ \Delta\mathbf{T} \end{bmatrix}_t = \begin{bmatrix} \boldsymbol{\Psi}_{\mathbf{XX}} & \boldsymbol{\Psi}_{\mathbf{XT}} \\ \boldsymbol{\Psi}_{\mathbf{TX}} & \boldsymbol{\Psi}_{\mathbf{TT}} \end{bmatrix} \begin{bmatrix} \mathbf{X} \\ \mathbf{T} \end{bmatrix}_{t-1} + \begin{bmatrix} \mathbf{H}_{\mathbf{X}} \\ \mathbf{H}_{\mathbf{T}} \end{bmatrix}_t = \boldsymbol{\Psi}\boldsymbol{\xi}_{t-1} + \mathbf{H}_t,$$

where the submatrix of parameters $\boldsymbol{\Psi}_{\mathbf{XX}}$ is a full rank $p \times p$, while $\boldsymbol{\Psi}_{\mathbf{XT}}$ is $p \times q$, $\boldsymbol{\Psi}_{\mathbf{TX}}$ is $q \times p$, and $\boldsymbol{\Psi}_{\mathbf{TT}}$ is $q \times q$.

Because **X** is the vector of *ordinary* variables, $\boldsymbol{\Psi}_{\mathbf{XX}}$ is full rank and the eigenvalues of $\mathbf{I}_p + \boldsymbol{\Psi}_{\mathbf{XX}}$ must be less one in absolute value. If the variables in **T** are the actual I(1) fundamental trends, as opposed to ordinary variables that serve as the conduits of the fundamental trends into the observable system, they must, as shown in Section 3.4, be mutually causally independent, requiring $\boldsymbol{\Psi}_{\mathbf{TT}} = \mathbf{0}_{qq}$, and *strongly exogenous*, requiring $\boldsymbol{\Psi}_{\mathbf{TX}} = \mathbf{0}_{qp}$ (Johansen 1995, p. 77; Juselius 2006, p. 263).

By analogy with the example in Section 3.3, the long-run causal structure of the ordinary variables can be defined as follows: Let **D** be the $p \times p$ matrix with the values of the main diagonal of $\boldsymbol{\Psi}_{\mathbf{XX}}$ on its main diagonal and zeroes elsewhere. Then,

(24) $$\mathbf{X}_t^{\infty} = -(\mathbf{D}^{-1}\boldsymbol{\Psi}_{\mathbf{XX}} - \mathbf{I})\mathbf{X}_t^{\infty} - \mathbf{D}^{-1}\boldsymbol{\Psi}_{\mathbf{XT}}\mathbf{T}_t^{\infty}.$$ [13]

The $\boldsymbol{\Psi}$ matrix in (22) can be decomposed analogously to the $\boldsymbol{\Pi}$ matrix in (2) such that $\boldsymbol{\Psi} = \boldsymbol{\alpha\beta}'$, where $\boldsymbol{\alpha}$ is $(p + q) \times r$ and $\boldsymbol{\beta}'$ is $r \times (p + q)$. The transitional causal structure embedded in $\boldsymbol{\Psi}$ that governs the transmission of shocks and ultimately determines the long-run causal structure reflected in (24) can be represented in this $\boldsymbol{\alpha\beta}'$-decomposition in the following *canonical* way: variables that are both cointegrated and directly causally connected are represented by the individual cointegrating relations expressed in $\boldsymbol{\beta}$ and the effects of causes are indicated by non-zero coefficients in $\boldsymbol{\alpha}$. To take a concrete example, consider a specific causal structure embedded in a CVAR like (22) and represented graphically in Figure 3. (With causal time-series graphs, we suppose that the arrows correspond to a one-period lag between a direct cause and its effect.) Thus, the causally canonical representation of Figure 3 would be given as

[13] Again, as noted in Section 3.3 with respect to equations (8) and (9) or (10) and (11), equation (24) is not a reduced form or long-run solution; it is the long-run causal structure. The matrix $\mathbf{D}^{-1}$ is simply a matrix of normalizing factors, enforcing the convention that long-run effects are placed on the left-hand side and causes on the right-hand side of the equation.

$$(25)\quad \Delta\boldsymbol{\xi}_t = \boldsymbol{\Psi}\boldsymbol{\xi}_{t-1} + \mathbf{H}_t = \begin{bmatrix} \psi_{11} & 0 & 0 & 0 & 0 & \psi_{16} & 0 \\ 0 & \psi_{22} & 0 & 0 & 0 & \psi_{26} & \psi_{17} \\ 0 & 0 & \psi_{33} & 0 & 0 & 0 & \psi_{37} \\ 0 & \psi_{42} & 0 & \psi_{44} & 0 & 0 & 0 \\ 0 & \psi_{52} & \psi_{53} & 0 & \psi_{55} & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} A \\ B \\ C \\ D \\ E \\ FT_1 \\ FT_2 \end{bmatrix}_{t-1} + \mathbf{H}_t$$

$$= \boldsymbol{\alpha}\boldsymbol{\beta}'\boldsymbol{\xi}_{t-1} + \mathbf{H}_t = \begin{bmatrix} \alpha_{11} & 0 & 0 & 0 & 0 \\ 0 & \alpha_{22} & 0 & 0 & 0 \\ 0 & 0 & \alpha_{33} & 0 & 0 \\ 0 & 0 & 0 & \alpha_{44} & 0 \\ 0 & 0 & 0 & 0 & \alpha_{55} \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & \beta_{16} & 0 \\ 0 & 1 & 0 & 0 & 0 & \beta_{26} & \beta_{17} \\ 0 & 0 & 1 & 0 & 0 & 0 & \beta_{37} \\ 0 & \beta_{42} & 0 & 1 & 0 & 0 & 0 \\ 0 & \beta_{52} & \beta_{53} & 0 & 1 & 0 & 0 \end{bmatrix} \begin{bmatrix} A \\ B \\ C \\ D \\ E \\ FT_1 \\ FT_2 \end{bmatrix}_{t-1} + \mathbf{H}_t$$

The rules governing the translation of the Figure 3 or any graph into the CVAR are straightforward:

(i) each single-variable direct causal pair or each collider is represented by a cointegrating relationship corresponding to a unique row of the $\boldsymbol{\beta}'$ matrix where the value of the parameter for the effect is normalized to unity;

(ii) there are as many adjustment parameters in $\boldsymbol{\alpha}$ as there are rows in $\boldsymbol{\beta}'$ (at most one per row) with the column of each non-zero parameter in $\boldsymbol{\alpha}$ corresponding to the row of one of the effects (i.e., corresponding to the row in which that variable is normalized to unity) in $\boldsymbol{\beta}'$;

(iii) if any variable is a cause, but not an effect with respect to all the other variables, it corresponds to a zero row in $\boldsymbol{\alpha}$ (and, thus, is weakly exogenous).

The $\boldsymbol{\beta}$ matrix thus tells us which variables are related causally and, therefore, connected by edges, and the $\boldsymbol{\alpha}$ matrix (equivalently the normalization of $\boldsymbol{\beta}'$) tells us which way the arrows point for those edges.

Except for trivial reorderings of the variables and rescalings, the CVAR (25) uniquely represents the causal graph in Figure 3. Algebraically, however, the matrices $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ are not unique. They can rotated to form other pairs ($\boldsymbol{\alpha}^*$ and $\boldsymbol{\beta}^*$) such that $\boldsymbol{\Psi} = \boldsymbol{\alpha}^*\boldsymbol{\beta}^{*\prime}$. The $\boldsymbol{\alpha}\boldsymbol{\beta}'$-representation and the $\boldsymbol{\alpha}^*\boldsymbol{\beta}^{*\prime}$- representation yield the same value of the likelihood function. The problem of causal search is to find empirical information other than the value of the likelihood function that would allow us to select the canonical representation as in CVAR (25) that corresponds to the graph of the data-generating process.

4.2. Trends, Cointegration, and Weak Exogeneity

4.2.1. *Formation and Sharing of Local Trends*

We can think of the causal graph of a system of I(1) variables as representing the channels of transmission of these trends. Each collider corresponds to the creation of a local trend, and the causal variables involved in the collider are cointegrated with the effect variable. The transmission of a local trend from one variable to a single other variable also implies the cointegration of the cause and the effect.

Although causal connections produce cointegration, cointegration itself is not essentially a causal notion. Instead, cointegration results either a) when a local trend is shared by two variables or b) whenever the number of variables sharing the same fundamental trends, whether or not they share the same local trends (i.e., whether or not they share the fundamental trends in the same proportions), exceeds the number of fundamental trends. Thus, in case b), if there is a set of variables each of which is driven by the same q fundamental trends, then any q+1 of them will be cointegrated. A causal connection is, thus, sufficient for the cointegration of the complete set of causes with their effect, but it is not necessary.

Proposition 7. Causal Cointegration: *If each member of the set of parents of a variable* C *in a causal graph is I(1), then the set of variables consisting of* C *and its parents, is cointegrated.*

It is convenient to write the fact that a set of variables is cointegrated as CI($\mathcal{Z}$), where $\mathcal{Z}$ is a set of variables with two or more members. Thus, if the variables A and B are cointegrated, we can write this as CI({A, B}). Two terms will prove useful:

Definition 9. *A* cointegrating group *is a set of variables in which every pair of variables shares the same common local trend – i.e., every pair is cointegrated*.

Definition 10. *A* collider group *is a set of variables that are cointegrated because they form a collider*.

The variables in a cointegration group *share* a single common local trend; while the variables in a collider group *generate* a new local trend. The same variable may be part of both a cointegration group and a collider group. Other sets of cointegrating variables may be in neither type of group. Davidson (1998, p. 91) introduces a useful concept, which we define here slightly differently that he does:

Definition 11. *A set of variables is* irreducibly cointegrating *(notated IC(·)) if, and only if, it does not contain a subset that is itself cointegrated.*

4.2.2. *A State-space Analysis of the CVAR*

It will prove useful to examine the relationship between weak exogeneity and the causal graph. Weak exogeneity is not in itself a causal property; rather it is a property related to the manner in which a likelihood function can be decomposed into a conditional and marginal probability distribution under a given parameterization (Engle, Hendry, and Richard 1983). Weak exogeneity is essentially the condition that guarantees that the parameters of interest can be estimated efficiently. Engle *et al*. define the term "parameter" differently from the way that we have defined it. In Section 4, we defined it as a structural feature of the DGP; while they define it as a feature of the particular model of the DGP (or some subpart of the DGP), so that a parameter in Engle *et al*.'s sense may, coincide with a parameter in our sense, but, equally, it may be a function of parameters in our sense.. To avoid confusion, we will use the term "coefficient" to mean parameter in Engle *et al.*'s sense (see Hoover 2001, pp. 175-177).

Given a DGP, the weak exogeneity status of its variables will depend on the model we estimate. So, for example, if (25) were the DGP and we estimated a model with precisely the form of the DGP, then FT_1 and FT_2 would be weakly exogenous in the conditional model $\{A, B, C, D, E\}_t | [\{A, B, C, D, E\}_{t-1}, \{FT_1, FT_2\}_{t-1}]$ for the coefficients ψ_{ij} or $(\alpha_{ji}, \beta_{ij})$, $i = 1, 2 . . , 5$, $j = 1, 2, . . ., 7$. Our main interest, however, will be in the case in which only a subset of the data of the DGP is modeled – leaving other variables latent. So, for example, we might consider data generated by (25), but observe only *B, C*, and *E*. These variables can be modeled in a CVAR form, but the coefficients of the model will not in general be the same as those of (25), though we could compute them if we knew the DGP. Still, we can ask the question whether we can decompose the likelihood function in a manner that renders some of the variables weakly exogenous with respect to the coefficients of a conditional model for the others.

We can notate this weak exogeneity using a new symbol "$\mapsto$", which means "is weakly exogenous for" and is to be distinguished from "$\rightarrow$", which means "directly causes." Thus, $\mathcal{X} \mapsto \mathcal{Y}$ can be read as "the variables in the set $\mathcal{X}$ are weakly exogenous for the coefficients of a CVAR model of $\mathcal{Y}$ conditional on $\mathcal{X}$." Since our main interest is the particular partition of the variables between the conditioned variables ($\mathcal{Y}$) and the conditioning variable ($\mathcal{X}$) and since the phrase in quotation marks in the last sentence is awkward, we will describe the situation indicated by $\mathcal{X} \mapsto \mathcal{Y}$ as "$\mathcal{X}$ is weakly exogenous for $\mathcal{Y}$," leaving the relativity of the weak exogeneity relationship to a particular set of coefficients embedded in a particular model implicit. If we know the causal graph of the DGP, then we can read the various weak exogeneity relationships for models of different subsets of variables from information in the causal graph. As a result, if we can identify weak exogeneity relationships for different subsets, we may be able to work backwards to determine which causal graphs could have generated them, in much the same way that graphical search models have typically worked backwards to determine a class of graphs consistent with facts of probabilistic independence.[14]

The object of the analysis is to use tests of long-run weak exogeneity in CVARs of the form of equation (1) *for observable variables only* to discover restrictions on allowable causal ordering of the underlying DGP (23). Long-run weak exogeneity corresponds to a zero row in the $\boldsymbol{\alpha}$ matrix of the CVAR, so a critical goal is, given a particular DGP, to determine what $\boldsymbol{\alpha}$ matrix it implies for a CVAR of the subset of observable variables.

Fundamental trends are assumed to be latent. In order to analyze cases in which some subsets of ordinary variables could be latent, partition $\mathbf{X}_t = [\mathbf{X}_{1t}, \mathbf{X}_{2t}]$, where the $\mathbf{X}_{1t}$ are observed and the $\mathbf{X}_{2t}$ (are treated as) unobserved. Then, rather than partitioning $\boldsymbol{\Psi}$ as in (23), partition it as

[14] Connecting weak exogeneity and, therefore, efficient estimation to causal order is reminiscent of LeRoy's (1995) approach to causality (cf. Hoover 2001, pp. 170-174). An importance difference, however, is that LeRoy defines causal orderings in terms of efficient estimation, while we seek only the implications of causal orderings defined in terms of structural parameterizations for weak exogeneity and tests of weak exogeneity as evidence of what the underlying structural relations might be.

$\boldsymbol{\Psi} = \begin{bmatrix} \overset{p\times p}{\mathbf{M}} & \overset{p\times m}{\mathbf{C}} \\ \underset{m\times p}{\mathbf{N}} & \underset{m\times m}{\mathbf{0}} \end{bmatrix}$, where the submatrices of parameters may or may not coincide with the $\boldsymbol{\Psi}_{\mathbf{IJ}}$,

depending on whether any ordinary variables are latent. The submatrix $\overset{p\times p}{\mathbf{M}} = \begin{bmatrix} \overset{p_1\times p_1}{\mathbf{M}_{11}} & \overset{p_1\times p_2}{\mathbf{M}_{12}} \\ \underset{p_2\times p_1}{\mathbf{M}_{21}} & \underset{p_2\times p_2}{\mathbf{M}_{22}} \end{bmatrix}$

contains the parameters of the ordinary variables. Only the parameters in $\mathbf{M}_{11}$ correspond to the p_1 observed ordinary variables, while the other $\mathbf{M}_{ij}$ contain parameters that correspond to the p_2 latent ordinary variables and m latent fundamental trends. The submatrix $\overset{p\times m}{\mathbf{C}} = \begin{bmatrix} \overset{m\times p_1}{\mathbf{C}_1'} & \overset{m\times p_2}{\mathbf{C}_2'} \end{bmatrix}'$

contains the parameters in $\mathbf{C}_1'$ corresponding to the effects of the latent fundamental trends on the observed ordinary variables and those in $\mathbf{C}_2'$ corresponding to the their effects on the unobserved ordinary variables. We suppose that $\mathbf{N} = \mathbf{0}$, which corresponds to the assumption that $\mathbf{T}_t$ are strongly exogenous.

A state-space representation of CVAR (23) can then be given:

$$\Delta\mathbf{X}_{1t+1} = \mathbf{M}_{11}\mathbf{X}_{1t} + \mathbf{M}_{12}\mathbf{X}_{2t} + \mathbf{C}_1\mathbf{T}_t + \boldsymbol{\varepsilon}_{1t+1}, \tag{26}$$

$$\Delta\mathbf{X}_{2t+1} = \mathbf{M}_{21}\mathbf{X}_{1t} + \mathbf{M}_{22}\mathbf{X}_{2t} + \mathbf{C}_2\mathbf{T}_t + \boldsymbol{\varepsilon}_{2t+1}, \tag{27}$$

$$\Delta\mathbf{T}_{t+1} = \boldsymbol{\eta}_{t+1}, \tag{28}$$

where $t = 0, 1, \ldots, n-1$, and $\mathbf{T}_0 = \mathbf{0}$ and $\mathbf{X}_0 = 0$. The shocks are partitioned into those affecting ordinary variables ($\boldsymbol{\varepsilon}$) and those affecting the latent variables ($\boldsymbol{\eta}$), with ($\boldsymbol{\varepsilon}_t$, $\boldsymbol{\eta}_t$) ~ i.i.d. $N_{p+m}(0, \boldsymbol{\Omega})$ and $\boldsymbol{\Omega} = \begin{pmatrix} \boldsymbol{\Omega}_\varepsilon & 0 \\ 0 & \boldsymbol{\Omega}_\eta \end{pmatrix}$. In keeping with the distinction between ordinary variables and fundamental trends, we assume that the eigenvalues of ($\mathbf{I}_p + \mathbf{M}$) are less than one in absolute value, so that the source of the nonstationarity of $\mathbf{X}_t$ is the fundamental trends, rather than its own dynamics.

The matrix **C** represents the proportions of fundamental trends present in observable variables but transmitted to them through latent causal connections and not via causal relationships among the observable variables. Thus, while the non-zero entries of **M** correspond to the edges in a causal graph, **C** is not given a direct graphical interpretation. The fundamental trends are embedded in **T**, but the variables included in **T** should be regarded as *local* trends, which may either be latent fundamental trends directly causing the observed variables or ordinary variables that carry some linear combination of fundamental trends and cause the observable variables. Therefore, $\boldsymbol{\Omega}_\eta > 0$ need not be diagonal.

Suppose that the DGP graph is described as in system (26)-(28), and we wish to know whether any of the *observed* variables ($\mathbf{X}_{1t}$) are weakly exogenous in a CVAR of the observed variables only. This comes down to the question of whether $\boldsymbol{\alpha}$ in that CVAR has any zero rows.

Define $\mathbf{C}^* = [\mathbf{M}_{12}\ \ \mathbf{C}_1]$ – i.e., as the matrix of the parameters governing the effects of latent ordinary variables and fundamental trends on observable ordinary variables. Let $s = \text{rank}(\mathbf{C}^*)$. Johansen's (2018) Theorem 1 shows that a solution to the system (26)-(28) requires that $s = m$; and, if $m < p_1$, then the observable variables cointegrate, whereas if $s = m$, they do not. Furthermore, if there is a zero row in $\mathbf{C}_\perp^*$, then there is a zero row in $\boldsymbol{\alpha}$ corresponding to the same variable, imply that the variable is weakly exogenous.

4.2.3. *Weak Exogeneity and Causal Order*

The state-space representation and Theorem 1 offer a tool for analyzing weak exogeneity for subsets of variables in the DGP. These, in turn, correspond in systematic ways to facts about the causal structure of the DGP itself. Consider some illustrative cases:

Case 1. Consider the causal graph in Figure 4 and assume that only the fundamental trends are unobserved, then the state-space analysis is given by

$$\mathbf{X}_t = \begin{bmatrix} A \\ B \\ C \end{bmatrix}_t,\ \mathbf{M} = \begin{bmatrix} \psi_{11} & 0 & 0 \\ 0 & \psi_{22} & 0 \\ \psi_{31} & \psi_{32} & 0 \end{bmatrix}, \text{ and } \mathbf{C} = \begin{bmatrix} \psi_{14} & \psi_{15} \\ \psi_{24} & \psi_{25} \\ 0 & 0 \end{bmatrix}.$$

so that $\mathbf{C}^* = \mathbf{C}$, $s = m = 2 < 3 = p_1 = p$ and

$$\mathbf{C}_\perp^* = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}.$$

The two zero rows in $\mathbf{C}_\perp^*$ imply that A and B are weakly exogenous for C (i.e., $\{A, B\} \mapsto C$). Notice that if the DGP were modified to add an edge between A and B (i.e., either $\psi_{23} \neq 0$ or $\psi_{32} \neq 0$), the analysis would be unaffected, since it relies on the $\mathbf{C}$ matrix only.

Case 2. Unfortunately, the simple mapping between weak exogeneity and causal connection suggested by Example 1 does not hold up. Consider Figure 5, which adds the variable D and edges connecting to other variables in Figure 4. Here there two fundamental trends, but three variables are parents in a collider at C. The analysis proceeds just as in Case 1 with the state-space formulation (omitting $\mathbf{M}$ as irrelevant) given by

$$\mathbf{X}_t = \begin{bmatrix} A \\ B \\ C \\ D \end{bmatrix}_t \text{ and } \mathbf{C}^* = \begin{bmatrix} \psi_{15} & \psi_{16} \\ \psi_{25} & \psi_{26} \\ 0 & 0 \\ 0 & \psi_{46} \end{bmatrix}; \text{ so that } \mathbf{C}_\perp^* = \begin{bmatrix} 0 & * \\ 0 & * \\ 1 & 0 \\ 0 & * \end{bmatrix},$$

where the stars indicate a non-zero value that must be chosen to conform to constraints implied by non-zero values of the $\mathbf{C}$ matrix. There is no loss in this open specification since it is a matter of indifference which of the non-unique $\mathbf{C}_\perp^*$ we obtain, as each will have the same zero rows and it is only the existence or non-existence of zero rows that matter for determining weak exogeneity.

The variables *A, B, C, D* are cointegrated (CI({*A, B, C, D*})); but, so is each three-member subset of these variables, implying *not* IC({*A, B, C, D*}). This is a robust finding: the parents in a collider are weakly exogenous only when the colliding set is irreducibly cointegrated.

Case 3. It is tempting to think that we might consider an irreducible subset of the variables in Figure 5, such as {*A, B, C*} and find the same weak exogeneity relations as we did in Figure 4. That, however, does not work. In analyzing the subset, we are effectively treating *D* as an unobservable variable. The state-space representation for this reduced system gives

$$\mathbf{X}_{1t} = \begin{bmatrix} A \\ B \\ C \end{bmatrix}_t, \ \mathbf{X}_{2t} = [D]_t, \text{ and } \mathbf{C}^* = \begin{bmatrix} 0 & \psi_{15} & \psi_{16} \\ 0 & \psi_{25} & \psi_{25} \\ \psi_{34} & 0 & 0 \end{bmatrix}.$$

In this case, $s = 3 > m$ and **C*** is full rank, so that $\mathbf{C}_{\perp} = \mathbf{0}$, implying that none of the variables is weakly exogenous. In treating, *D* as unobservable, removes it from the graph while, at the same time, not eliminating the fact that it provides a conduit from the fundamental trends to *C* that is distinct from the observable conduits, *A* and *B.* It is as if the graph of Figure 5 has been transformed into Figure 6, where the dashed arrow indicates an causal connection between FT_2 and *C*, mediated by *D* in the DGP but not observable in the CVAR of the subset {*A*, *B*, *C*}.

Case 4. In Case 3, weak exogeneity failed to obtain, even though the causal connections were genuine. It can also happen that weak exogeneity does obtain, even when causal connections are missing. Consider Figure 7. Although the graph shows *not* ($A \rightarrow C$) and *not* ($B \rightarrow D$) and *not* ($B \rightarrow E$), although *B* directly causes *E*. Using the same state-space methods, however, we can show that $\{A, B\} \mapsto \{C, D, E\}$. And, looking at subsets of variables $\{A, B\} \mapsto D$. Thus, {*A*, *B*, *D*} have the same apparent pattern of weak exogeneity as found for {*A*, *B*, *C*}in Case 1 (Figure 4); yet these variables do not form a collider group in Figure 7. Again, we have a failure of *irreducible* cointegration: CI({*A*, *B*, *D*}), but also CI({*A*, *D*}).

Case 5. Weak exogeneity may fail to track *direct* cause. Consider a causal chain : $FT \rightarrow A \rightarrow B \rightarrow C$. Suppose $A \mapsto C$ but also that $B \mapsto C$ and {*A*, *B*, *C*} form a cointegration group, then $B \rightarrow C$ and *not* ($A \rightarrow C$). Similarly, if CI({*A*, *B*, *C*, *D*}) and $\{A, D\} \mapsto C$ and $\{B, D\} \mapsto C$, but {*A*, *B*} is a cointegration group with $A \mapsto B$, then {*B,C,D*} form a collider group (i.e., $B \rightarrow C \leftarrow D$) but {*A,C,D*} does not form a collider group (i.e., *not* ($A \rightarrow C \leftarrow D$). In addition to weak exogeneity and irreducible cointegration, a direct cause must be adjacent to its effect. One lesson is that, if every member of the same cointegration group are interchangeable as one of the potential parents in a collider group, then any one of them that is weakly exogenous for the other member cannot be the direct cause. By similar reasoning, if members of the same cointegration group are interchangeable as the potential child in a collider group, then the one that is weakly exogenous for the others is the only one that can be the actual direct effect of the parents.

These cases show us how to read weak exogeneity off a causal graph. There are two cases:

1. Within a set of variables that form a cointegration group, a particular variable is weakly exogenous for the group if, and only if, it the sole source of the local trend that cointegrates the group.

2. The parents in any set of variables that form a collider group in which two or more local trends are combined are weakly exogenous for the child in the collider group, provided that the number of variables in the group is fewer than 1 + the number of fundamental trends carried by those variables. Similarly, if such a collider exists in the graph, then in any set that replaces one or more weakly exogenous parents with a variable in the same cointegration group as that parent, provided the variable is itself weakly exogenous for the parent, will also be weakly exogenous for the child. And any variable that is weakly exogenous for the child either as a parent or as a member of the same cointegration group that replaces the parent will be weakly exogenous for a variable that replaces the child from a cointegration group that includes the child. (Thus, in Figure 3, $\{FT_1, FT_2\} \mapsto B$, but in the set that replaces *B* with *D*, which are both in the same cointegration group, $\{FT_1, FT_2\} \mapsto D$. And, in Figure 7, in the collider $\{A, C, E\}$, $\{A, C\} \mapsto E$; but in the set in which *B* replaces *C* (both in the same collider group), $\{A, B\} \mapsto E$.)

The inferential lessons of Cases 1-4, can be summarized in three rules, consistent with visual reading of the graph :

Rule 1. *If* $\mathrm{A} \mapsto \mathrm{B}$, *then* not $\mathrm{B} \rightarrow \mathrm{A}$.
Rule 1 simply says that causation cannot run against the direction of weak exogeneity.

Rule 2. *In a cointegration group, if* $\mathrm{A} \mapsto \mathrm{B}$ *and there is no* C *such that* $\mathrm{A} \mapsto \mathrm{C}$ *and* $\mathrm{C} \mapsto \mathrm{B}$*, then* $\mathrm{A} \rightarrow \mathrm{B}$.
Rule 2 says that bivariate weak exogeneity coincides with direct causation, provided that the variables are adjacent.

Rule 3. *A triple of variables forms a collider* $\mathrm{A} \rightarrow \mathrm{C} \leftarrow \mathrm{B}$*, if i) IC(*$\{\mathrm{A}, \mathrm{B}, \mathrm{C}\}$*); ii)* $\{\mathrm{A}, \mathrm{B}\} \mapsto \mathrm{C}$*; iii) it is not the case that* A *is a member of a cointegration group* $\mathcal{Z}$ *such that, for any member* $\mathrm{D} \in \mathcal{Z}$ *(excluding* A*),* $\mathrm{A} \mapsto \mathrm{D}$ *and* $\{\mathrm{B}, \mathrm{D}\} \mapsto \mathrm{C}$*, and mutatis mutandis for* B*; and iv) it is not the case that* C *is a member of a cointegration group* $\mathcal{Z}$ *such that for any member* $\mathrm{D} \in \mathcal{Z}$ *(excluding* C*) that* $\mathrm{D} \mapsto \mathrm{C}$.
Rules 3 says that if two variables are weakly exogenous for a third, they form a triple, provided that each of the weakly exogenous variables is adjacent to the third variable (established by conditions iii) and iv)).

4.3. Long-run Causal Search in a Causally Sufficient Graph

Davidson (1998, section 3) proposes a search algorithm that identifies every irreducible cointegrating set of variables within a CVAR. He then uses that information where possible to identify the cointegrating relations in the $\boldsymbol{\beta}'$ matrix. This strategy is successful in some cases and not others. There is an analogy with causal search for stationary variables. Despite the slogan, "correlation is not causation," it is sometimes possible to infer causal direction from tests of *unconditional* dependence. For example, for a causally sufficient set of three variables with an acyclical data-generating process, if *A* and *C* are not correlated, but *A* and *B* and *B* and *C* are correlated, then $A \rightarrow B \leftarrow C$ is the only consistent causal graph. In most cases, however, unconditional independence is not enough. Relations of *conditional* dependence and independence provides a richer source of information for inferring the *direction*, as well as the existence of causal edges (see Section 2.2 above). Davidson's schema places cointegration in something like the role of unconditional independence (or correlation) in the stationary case.

Davidson's inferential scheme is less informative than it could be, in part, because it fails to acknowledge the role of fundamental trends (either explicit or latent) and, in part, because it fails to exploit all the available evidence on causal asymmetry. In addition to Davidson's assessment of the irreducible cointegration for every subset of the variables, complete assessment of weak exogeneity among those irreducible cointegrating subsets can function in something like the role of conditional independence, when processed according to the three rules in the last section and may provide richer, empirically grounded information about the identification of the CVAR. As with causal search in the stationary case, the application of these rules will not identify every possible causal graph, but will sometimes be able to partially or completely uncover the underlying causal structure.

Consider the DGP in Figure 3 and assume that its variables are causally sufficient and all (including the fundamental trends) are observed. We are interested in the logic of causal inference rather than the statistical problem of inference, so we also assume that we know the correct facts with respect to the cointegration rank and cointegration and weak exogeneity among any subset of variables. (In the language of the causal search literature, we assume that we have an *oracle*.) Can we use this information to recover the graph of the DGP?

The inference problem can be viewed as how to place the zero and non-zero coefficients the $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}'$ matrices in equation (25).

Given that we know that the cointegration rank is 5, we know that there are two fundamental trends. This implies that $\boldsymbol{\alpha}$ is 7×5 and $\boldsymbol{\beta}'$ 5×7. Since FT_1 and FT_2 are weakly exogenous with respect to all other variables in the system, we may conclude that, even if they are not identical with the fundamental trends (which in this case, of course, they are), they at least are the unique sources introducing those trends into the system. And we are entitled to enter zeroes in the entire rows of $\boldsymbol{\alpha}$ corresponding to FT_1 and FT_2. Without loss of generality, we may enter non-zero α_{ij} along the main diagonal of the submatrix excluding the FT_1- and FT_2-rows of $\boldsymbol{\alpha}$ and zeroes everywhere else. Similarly, we may enter ones on the main diagonal of the submatrix of $\boldsymbol{\beta}'$ that excludes the last two columns.

With two fundamental trends, no irreducible cointegrating relation can involve more than three variables. Exhaustive consideration along Davidson's lines would produce 21 possible cointegrating pairs and 35 possible cointegrating triples. Similarly, we need to consider possible weak exogeneity of variables within each irreducibly cointegrating subset. Most of subsets are not irreducibly cointegrating or do not contain weakly exogenous variables, so rather than listing all the subsets systematically, we just note the salient ones.

From the facts that CI($\{A, FT_1\}$) and that there are no other variables in this cointegration group and that $FT_1 \mapsto A$, Rule 2 implies $FT_1 \rightarrow A$, which justifies the placement of β_{16} in row 1 of $\boldsymbol{\beta}'$ and zeroes in the remaining unassigned places in that row. Analogous reasoning with respect to $\{C, FT_2\}$ implies $FT_2 \rightarrow C$ and justifies the placement of β_{37} and the zeroes in row 3. and again with respect to $\{B, D\}$, analogous reasoning justifies the placement of β_{42} and the zeroes in row 4. In addition, in this case, Rule 1 and the fact that $B \mapsto D$ imply that *not* ($D \rightarrow B$) and justify the zero in row 2, column 4.

Rule 3 and the facts that IC($\{FT_1, FT_2, B\}$, that B is not part of a cointegration group with either FT_1 or FT_2, and that $\{FT_1, FT_2\} \mapsto B$ allows us to identify the collider

$FT_1 \rightarrow B \leftarrow FT_2$ and justifies the placement of β_{26} and β_{27} and the remaining zeroes in row 2 of $\boldsymbol{\beta}'$.

Rules 3 and the facts that IC({*B*, *C*, *E*}), ({*B*, *C*} $\mapsto E$, and *not* ($C \mapsto FT_2$), with which it forms a cointegration group, allows us to identify the collider $B \rightarrow E \leftarrow C$ and justifies the placement of β_{52} and β_{53} and the zeroes in row 5 of $\boldsymbol{\beta}'$.

With that we were able to recover the entire DGP graph using only the facts of cointegration and weak exogeneity.

4.4. LONG-RUN CAUSAL SEARCH IN THE PRESENCE OF LATENT TRENDS

The CVARs typically estimated in practice most often do not contain variables that are weakly exogenous for the whole system, which could, therefore, be identified as the conduit of the fundamental trends to the other variables in the system. It is, therefore, worth considering how the principles of search might operate when fundamental trends are latent variables. It is possible to apply the rules of Section 4.2 the variables generated according to equation (25) and to treat only the ordinary variables (*A*, *B*, *C*, *D*, *E*) as observed and the fundamental trends (FT_1 and FT_2) as unobserved. For some of the causal edges, the reasoning of Section 4.3 is still applicable, and we would be able to infer the edges shown in Figure 8: $B \rightarrow A$ and $B \rightarrow E \leftarrow C$. The remainder of Figure 8 requires further comment.

We are unable to infer the edges between FT_1, FT_2 and *A*, *B*, and *C* for the simple reason that the two fundamental trends are not observed and the inference of the edges in which they are involved requires their observability. However, we do know from the fact that the cointegration rank is 5 that there are two fundamental trends. What we cannot say, however, is that those two trends enter directly into the observable system. They may, in fact, be transmitted through ordinary variables that are also latent. We do, know, however, that must enter through *A*, *B*, or *C*. If that were not the case and a fundamental trend entered through *D* or *E*, we would not have found that CI({*B*, *D*}) or {*B*, *C*} $\mapsto E$. This is indicated in Figure 8 by the oval enclosing the ordinary variables and the circles (indicating their latency) around the fundamental trends. The arrows running from the latent fundamental trends to the oval, stopping short of the particular variables indicates that we know that these variables are caused by these trends, albeit we do not exactly what the connections are. Thus, instead of (25), we can fill in the causally ordered CVAR equation with the ambiguous information depicted in Figure 8, where the question marks indicate parameters that correspond to possible, but yet to be determined causal edges:

(29) $\Delta \boldsymbol{\xi}_t = \boldsymbol{\Psi}\boldsymbol{\xi}_{t-1} + \mathbf{H}_t$

$$= \boldsymbol{\alpha}\boldsymbol{\beta}'\boldsymbol{\xi}_{t-1} + \mathbf{H}_t = \begin{bmatrix} \alpha_{11} & 0 & 0 & 0 & 0 \\ 0 & \alpha_{22} & 0 & 0 & 0 \\ 0 & 0 & \alpha_{33} & 0 & 0 \\ 0 & 0 & 0 & \alpha_{44} & 0 \\ 0 & 0 & 0 & 0 & \alpha_{55} \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} 1 & ? & ? & 0 & 0 & ? & ? \\ ? & 1 & ? & 0 & 0 & ? & ? \\ ? & ? & 1 & 0 & 0 & ? & ? \\ 0 & \beta_{42} & 0 & 1 & 0 & 0 & 0 \\ 0 & \beta_{52} & \beta_{53} & 0 & 1 & 0 & 0 \end{bmatrix} \begin{bmatrix} A \\ B \\ C \\ D \\ E \\ L_1 \\ L_2 \end{bmatrix}_{t-1} + \mathbf{H}_t$$

Here the two variables, L_1 and L_2 are latent variable that may either be the underlying fundamental trends or ordinary variables that transmit distinct linear combinations of the fundamental trends to the system. In the later case, the covariance between their error terms in $\boldsymbol{\Omega}$ may be nonzero.

Neither the graph nor (29) conveys all the information that we have. We know, for instance, that there are two fundamental trends and that at least one of the fundamental trends must be a causal influence on each of *A, B,* and *C*. If that were not so, then the only way that all three variables could carry the trends and be irreducibly cointegrated would be for them to form a collider group in which one pair is weakly exogenous for the remaining variable. Given the DGP, we know that the weak exogeneity search would not have found that. Furthermore, we know that no two of *A, B,* and *C* can have a common latent cause. If that were not true, that pair would form a cointegration group, which the search for cointegrating pairs would not have, in fact, discovered. These two conclusions imply that each of the three observed variables carries the fundamental trends in distinct proportions. These facts place restrictions on how the last two columns of the $\boldsymbol{\beta}'$ in (29) can be filled in to be consistent with the DGP.

5. Conclusion

In the history of econometrics, the problem of identification and the notion of causal order have long been connected – both in the work of Simon and the early Cowles Commission program and in the literature on SVARs. Typically, economists have relied heavily on the idea that *a priori* restrictions derived somehow from economic theory would provide the needed identification. Recent work on graphical causal modeling, however, has shown that there is often unexploited information that could provide a firmer, empirical basis for identification. In the case of cross-sectional data or the contemporaneous causal orderings of SVARs, the graphical causal modelers have stressed the information contained in conditional independence relationship encoded in the probability distribution of the data. Conditional independence may also be a resource in the case of the long-run dynamics of the CVAR, although the fact that nonstationary data involves non-standard distributions poses some challenges. We have suggest here that nonstationary data also present the opportunity to take a different approach.

CVARs are nearly decomposable systems in Simon's sense: the long-run causal order of such systems emerges from complex short-run dynamics to reveal a simpler set of temporally aggregated relationships. Our focus has been on how that emergence happens and what causal relationships underwrite it. The fact that the short-run behavior of the system is largely distinct from the longer-run, trend driven behavior gives us a resource, not available for the cases of cross-sectional data or the contemporaneous causal order of the SVAR, of, in effect, tracing the flows of trends as they combine into new local trends, while, at the same, time preserving their essential distinct identities. Focus on the trends raises questions about the nature of trends.

Where do trend come from? We have argued that, while it is possible for systems of equations to generate trend-like behavior on the basis of particular fortuitous parameterizations, such "trends" are unlikely to be robust and would require a special economic explanation. More likely trends arise from economic processes that generate variables that display trend behavior essentially. Once a distinction is drawn between ordinary variables (stationary processors) and fundamental trends (nonstationary processors), it is clear that a more robust account for nonstationary behavior is that it is transmitted from its fundamental sources to variables that without these fundamental trends as direct or indirect causes would not naturally be nonstationary. In typical CVAR analysis, econometricians mostly do not find variables that themselves can be identified as the source of fundamental trends. This suggests that, in most cases, fundamental trends are latent variables, and any sort of structural or causal analysis of CVARs must account for their latency.

We suggested – somewhat informally – that combining Davidson's suggestion of a comprehensive search for sets of irreducible cointegrating relations with a similar comprehensive search of weak exogeneity among those sets could provide a non-*a priori* empirical basis for discovering identifying restrictions on cointegrating relations, as well as information on causal direction. We showed that in a simple example, the complete causal graph of the CVAR could be recovered. But, in most cases in the face of latent variables, these restrictions are unlikely to provide complete identification. Nevertheless, as in our illustration, some of the cointegrating relations may be identified, even when there are latent trends. It is also possible that, in some cases, it would be possible to recover estimates of the trends using state-space methods (see e.g., Johansen and Tabor 2017). Finally, viewing the CVAR through the lens of latent fundamental trends reinforces Juselius's advocacy of simple-to-general modeling in the CVAR context (Juselius 2006, ch. 22, esp. sections 22.2.3 and 22.3). Cointegrating relations are robust to widening the data set to include more variables. The aim of such widening can be seen as an effort to discover the observable variables that are the counterpart of the latent trends in narrower data sets.

References

Batterman, Robert W. (2001) *The Devil in the Details : Asymptotic Reasoning in Explanation, Reduction and Emergence.* Oxford: Oxford University Press.

Boumans, Marcel. (2005) *How Economists Model the World into Numbers*. London: Routledge.

Cooley, Thomas F. and Stephen F. LeRoy. (1985) "Atheoretical Macroeconomics: A Critique," *Journal of Monetary Economics* vol. 16(3), 283-308.

Cooper, Gregory F. (1999). 'An Overview of the Representation and Discovery of Causal Relationships using Bayesian Networks', in Clark Glymour and Gregory F. Cooper, editors. *Computation, Causation, and Discovery*, American Association for Artificial Intelligence, Menlo Park, CA and MIT Press, Cambridge, MA, pp. 3-64.

Davidson, James. (1998) "Structural Relations, Cointegration and Identification: Som Simple Results and Their Application," *Journal of Econometrics* 87(1), pp. 87-113.

Demiralp, Selva and Kevin D. Hoover. (2003) "Searching for the Causal Structure of a Vector Autoregression," *Oxford Bulletin of Economics and Statistics* 65(supplement):745-767.

Demiralp, Selva, Kevin D. Hoover, and Stephen J. Perez. (2008) "A Bootstrap Method for Identifying and Evaluating a Structural Vector Autoregression", *Oxford Bulletin of Economics and Statistics* 70(4), pp. 509-533.

Duarte, Pedro Garcia and Kevin D. Hoover (2012) "Observing Shocks," in Harro Maas and Mary Morgan, editors. *Histories of Observation in Economics, supplement to History of Political Economy 44(5),* 226-249.

Engle, Robert F., David F. Hendry, and Jean-François Richard. (1983) "Exogeneity," *Econometrica* 51(2), 277-304.

Granger, C.W.J. (1969) "Investigating Causal Relations by Econometric Models and Cross-spectral Methods," *Econometrica* 37(3), pp. 424-438.

Haavelmo, Trygve. (1944) "The Probability Approach in Econometrics," *Econometrica* 12(Supplement), pp. iii-115.

Hendry, David F. (1995) *Dynamic Econometrics.* Oxford: Oxford University Press.

Hood, William C. and Tjalling Koopmans, editors. (1953) *Studies in Econometric Method,* Cowles Commission Monograph vol. 14. New York: Wiley, pp. 1-26.

Hoover, Kevin D. (2001) *Causality in Macroeconomics.* Cambridge: Cambridge University Press.

Hoover, Kevin D. (2004) "Lost Causes," *Journal of the History of Economic Thought* 26(2), pp. 149-164.

Hoover, Kevin D., "On the Reception of Haavelmo's Econometric Thought," *Journal of the History of Economic Thought* 36(1), March 2014, 45-65.

Hoover, Kevin D. (2015) "The Ontological Status of Shocks and Trends in Macroeconomics," *Synthese*192(3), pp. 3509–3532

Hoover, Kevin D. and Òscar Jordà. (2001) "Measuring Systematic Monetary Policy," *Federal Reserve Bank of St. Louis Review* 83(4), 113-137.

Hoover, Kevin D. and Katarina Juselius. (2015) "Trygve Haavelmo's Experimental Methodology and Scenario Analysis in a Cointegrated Vector Autoregression," *Econometric Theory* 31(2), pp. 249-274.

Johansen, Søren. (1995) *Likelihood-based Inference in Cointegrated Vector Autoregressive Models.* Oxford: Oxford University Press.

Johansen, Søren. (2018) "Cointegration and Adjustment in the CVAR(∞) Representation of Some Partially Observed CVAR(1) Models," Department of Economics, University of Copenhagen, unpublished typescript, 1 May 2018.

Johansen, Søren and Katarina Juselius. (1994) "Identification of the Long-run and the Short-run Structure: An Application to the ISLM Model," *Journal of Econometrics* 63(1), pp. 7-36.

Johansen, Søren and Katarina Juselius. (2014) "An Asymptotic Invariance Property of the Common Trends Under Linear Transformation of the Data," *Journal of Econometrics* 178, pp. 310-315.

Johansen, Søren and Morten Nyboe Tabor. (2017) "Cointegration between Trends and Their Estimators in State Space Models and Cointegrated Vector Autoregressive Models," *Econometrics* 5(36), pp. 3-15.

Juselius, Katarina. (2006) *The Cointegrated VAR Model.* Oxford: Oxford University Press.

Koopmans, Tjalling C., editor. (1950) *Statistical Inference in Dynamic Economic Models*, Cowles Commission Monograph 10. New York: Wiley.

LeRoy, Stephen F. (1995) "Causal Orderings," in Kevin D. Hoover, editor. *Macroeconometrics: Developments, Tensions, and Prospects.* Boston: Kluwer, pp. 235-252.

Liu, Ta-Chung. (1960) "Underidentification, Structural Estimation, and Forecasting," *Econometrica* 28(4), pp. 855–865.

Morgan, Mary S. (1990) *The History of Econometric Ideas*. Cambridge: Cambridge University Press.

Pearl, Judea. 2000. *Causality: Models, Reasoning, and Inference*. Cambridge: Cambridge University Press.

Phiromswad, Piyachart and Kevin D. Hoover (2013) "Selecting Instrumental Variables: A Graph-Theoretic Approach," Duke University Economics Department, unpublished typescript.

Reichenbach, H. (1956) *The Direction of Time*, Berkeley: University of California Press.

Robert F. Engle and Clive W. J. Granger (1987) "Cointegration and Error Correction Representation: Estimation and Testing," *Econometrica* 55(2), pp. 251-276.

Simon, Herbert A. (1953) "Causal Order and Identifiability," in Hood and Koopmans (1953), pp. 49-74.

Simon, Herbert A. (1996) *The Sciences of the Artificial,* 3rd ed. Cambridge, MA: MIT Press.

Simon, Herbert A. and Albert Ando. (1961) "Aggregation of Variables in Dynamic Systems," *Econometrica* 29(2), 111-138.

Simon, Herbert A. and Yumi Iwasaki. (1988) "Causal Ordering, Comparative Statics, and Near Decomposability," *Journal of Econometrics* 39(1/2), 149-173.

Simon, Herbert A. and Nicholas Rescher. (1961) "Causes and Counterfactuals," *Philosophy of Science* 33(4), pp. 323-340.

Sims, Christopher A. (1980) "Macroeconomics and Reality," *Econometrica*, 48(1):1-48.

Sims, Christopher A. (1982) "Policy Analysis with Econometric Models," *Brookings Papers on Economic Activity*, 13(1):107-52.

Sims, Christopher A. (1986) "Are Forecasting Models Usable for Policy Analysis?" *Federal Reserve Bank of Minneapolis Quarterly Review*, 10:2-15.

Spirtes, Peter, Clark Glymour, and Richard Scheines. (2000) *Causation, Prediction, and* Search, 2nd edition. Cambridge, MA: MIT Press.

Swanson, Norman R. and Clive W.J. Granger. (1997). "Impulse Response Functions Based on a Causal Approach to Residual Orthogonalization in Vector Autoregressions," *Journal of the American Statistical Association* 92(437), pp. 357-367.

Wold, Hermann. (1960) "A Generalization of Causal Chain Models," *Econometrica* 28(2), pp. 277-326.

Figure 1

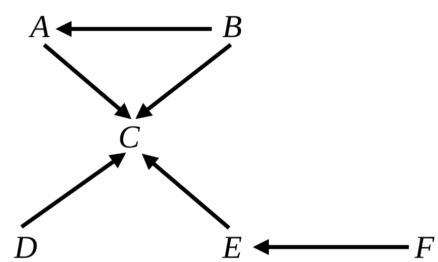

Figure 2.A

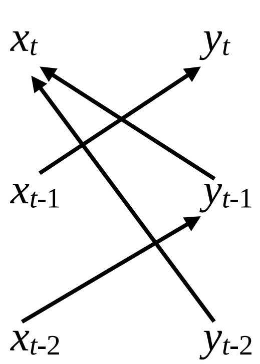

Figure 2.B

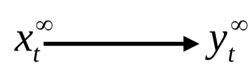

Figure 3

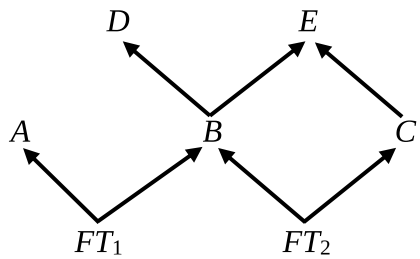

Figure 4

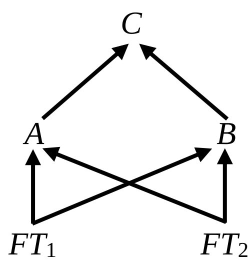

Figure 5

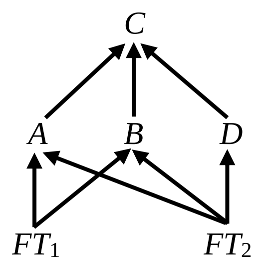

Figure 6

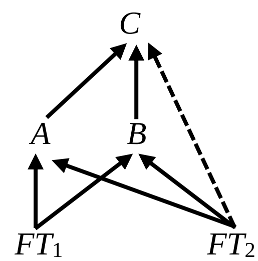

Figure 7

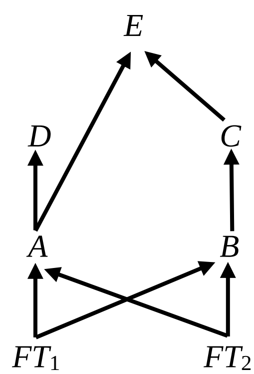

Figure 8